Geometry & Topology 21 (2017) 1583–1630

Genus-two trisections are standard

JEFFREY MEIER
ALEXANDER ZUPAN

We show that the only closed 4–manifolds admitting genus-two trisections are $S^2 \times S^2$ and connected sums of $S^1 \times S^3$, $\mathbb{CP}^2$ and $\overline{\mathbb{CP}}^2$ with two summands. Moreover, each of these manifolds admits a unique genus-two trisection up to diffeomorphism. The proof relies heavily on the combinatorics of genus-two Heegaard diagrams of S^3. As a corollary, we classify tunnel number one links with an integral cosmetic Dehn surgery.

57N12, 57R65; 57M25

1 Introduction

A trisection of a smooth 4–manifold X is a decomposition of X into three pieces, each of which is diffeomorphic to a 4–dimensional handlebody. Trisections have been introduced by Gay and Kirby [6] as a 4–dimensional analogue to Heegaard splittings of 3–manifolds. An explicit goal of the theory of trisections is to provide a vehicle to apply 3–manifold techniques to problems in 4–dimensional topology. In their seminal paper, they prove that every closed, orientable 4–manifold admits a trisection, and any two trisections for the same 4–manifold become isotopic after some number of iterations of a natural stabilization operation, akin to the Reidemeister–Singer theorem in dimension three.

While a Heegaard splitting of a 3–manifold has only one complexity parameter (the genus of the splitting), a trisection comes equipped with two such parameters. A (g,k)*–trisection* of a closed, orientable, smooth 4–manifold X is a decomposition

$$X = X_1 \cup X_2 \cup X_3,$$

where each X_i is a 4–dimensional 1–handlebody obtained by attaching k 1–handles to one 0–handle, each $Y_{ij} = X_i \cap X_j$ is a 3–dimensional genus-g handlebody, and $\Sigma = X_1 \cap X_2 \cap X_3$ is a closed genus-g surface, which we call a *trisection surface*. Observe that Σ is a Heegaard surface for each 3–manifold $\partial X_i = \#^k(S^1 \times S^2)$; hence, the definition implies that $g \geq k$.

Published: 10 May 2017 DOI: 10.2140/gt.2017.21.1583

Since the field is still in its infancy, there are many basic questions about trisections which have not yet been answered. In terms of 3–manifolds, the only manifold Y with a genus-zero Heegaard splitting is S^3, and the manifolds with genus-one splittings (lens spaces, $S^1 \times S^2$ and S^3) are neatly parametrized by the extended rational numbers. This leads to the following general question.

Question 1.1 To what extent can we enumerate (g,k)–trisections for small values of g?

In addition to proving the existence of trisections, Gay and Kirby relate them to handle decompositions. They show that a 4–manifold admitting a (g,k)–trisection has a handle decomposition with one 0–handle, k 1–handles, $(g-k)$ 2–handles, k 3–handles and one 4–handle. As such, it is easy to classify (g,g)–trisections. If X has a (g,g)–trisection, then X has a handle decomposition with no 2–handles; therefore, by Laudenbach and Poénaru [14], we have that X is diffeomorphic to $\#^g(S^1 \times S^3)$ (where $\#^0(S^1 \times S^3) = S^4$). It follows that there is a unique genus-zero trisection: the standard $(0,0)$–trisection of S^4 into three 4–balls, two of which are glued to each other along one hemisphere of their 3–sphere boundaries, and the third of which is attached along the resulting 3–sphere boundary.

Note that every Heegaard splitting $Y = H_1 \cup_\Sigma H_2$ can be presented by a Heegaard diagram (in fact, infinitely many Heegaard diagrams), which is a triple denoted (Σ,α,β). This notion has an analogue in terms of trisections: a *(g,k)–trisection diagram* is a quadruple $(\Sigma,\alpha,\beta,\gamma)$ such that each Heegaard diagram (Σ,α,β), (Σ,β,γ) and (Σ,α,γ) presents a genus-g splitting of $\#^k(S^1 \times S^2)$. Every trisection diagram yields a straightforward construction of a trisected 4–manifold X; see Section 2.

The manifolds that admit genus-one trisections are not difficult to classify; they are described in detail by Gay and Kirby [6], and diagrams of these trisections are pictured in Figure 1. There are precisely two manifolds which admit $(1,0)$–trisections: $\mathbb{CP}^2$ and $\overline{\mathbb{CP}}^2$. Thus, the situation is somewhat different here than in dimension three, suggesting another basic question.

Question 1.2 What is the smallest value of g for which there are infinitely many 4–manifolds X which admit a (g,k)–trisection for some k?

In [18], the authors prove that the answer to Question 1.2 satisfies $g \le 3$. It is shown that if X is the double cover of S^4 branched over a k–twist spun 2–bridge knot where $k \neq \pm 1$, then X admits a minimal genus $(3,1)$–trisection. Hence, Question 1.2 reduces to understanding genus-two trisections. This brings us to the main result of the present paper.

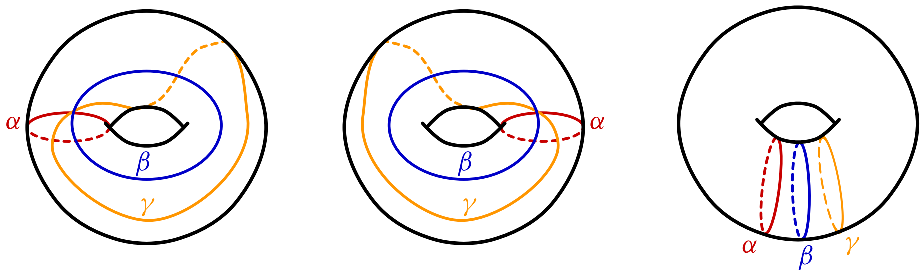

Figure 1: The three possible genus-one trisection diagrams: the (1,0)–trisection diagrams for $\mathbb{CP}^2$ (left) and $\overline{\mathbb{CP}}^2$ (middle), and the (1,1)–trisection diagram for $S^1 \times S^3$ (right)

Theorem 1.3 *If X admits a genus-two trisection, then X is either $S^2 \times S^2$ or a connected sum of $S^1 \times S^3$, $\mathbb{CP}^2$ and $\overline{\mathbb{CP}}^2$ with two summands. Moreover, each of these 4–manifolds has a unique genus-two trisection up to diffeomorphism.*

Proof Suppose X admits a $(2,k)$–trisection $X = X_1 \cup X_2 \cup X_3$. If $k = 2$, then trivially X is diffeomorphic to $(S^1 \times S^3) \# (S^1 \times S^3)$ and has a trisection diagram (α, β, γ) such that $\alpha = \beta = \gamma$. If $k = 1$, then by Proposition 7.3, the trisection is reducible and can be written as the connected sum of one of the standard $(1,0)$–trisections and the standard $(1,1)$–trisection shown in Figure 1. It follows that X is diffeomorphic to $(S^1 \times S^3) \# \mathbb{CP}^2$ or $(S^1 \times S^3) \# \overline{\mathbb{CP}}^2$. Finally, if $k = 0$, then by Theorem 6.5, the trisection has a diagram homeomorphic to one of the standard diagrams pictured in Figure 2, and thus X is diffeomorphic to $S^2 \times S^2$, $\mathbb{CP}^2 \# \mathbb{CP}^2$ or $\mathbb{CP}^2 \# \overline{\mathbb{CP}}^2$ by [6]. □

The bulk of this paper is devoted to proving Theorem 6.5, the classification of $(2,0)$–trisections, by understanding $(2,0)$–trisection diagrams $(\Sigma, \alpha, \beta, \gamma)$. In this case, each of (Σ, α, β), (Σ, β, γ) and (Σ, γ, α) is a genus-two Heegaard diagram for S^3, and these are well-studied objects in 3–manifold topology. A result of Homma, Ochiai and Takahashi [11] states that every nontrivial genus-two Heegaard diagram of S^3 can be reduced algorithmically (see Section 2), and this is the main input into our proof of Theorem 6.5.

The case of $(2,1)$–trisections, which is much less involved, is taken up in Section 7. In fact, the classification of $(2,1)$–trisections has been simplified by the authors and Schirmer [17] as part of a larger classification of *unbalanced* trisections (ie trisections in which the 4–dimensional 1–handlebodies are allowed to have different genus), including unbalanced genus-two trisections.

As a corollary to the main theorem, we obtain the following result; see Section 7 for definitions.

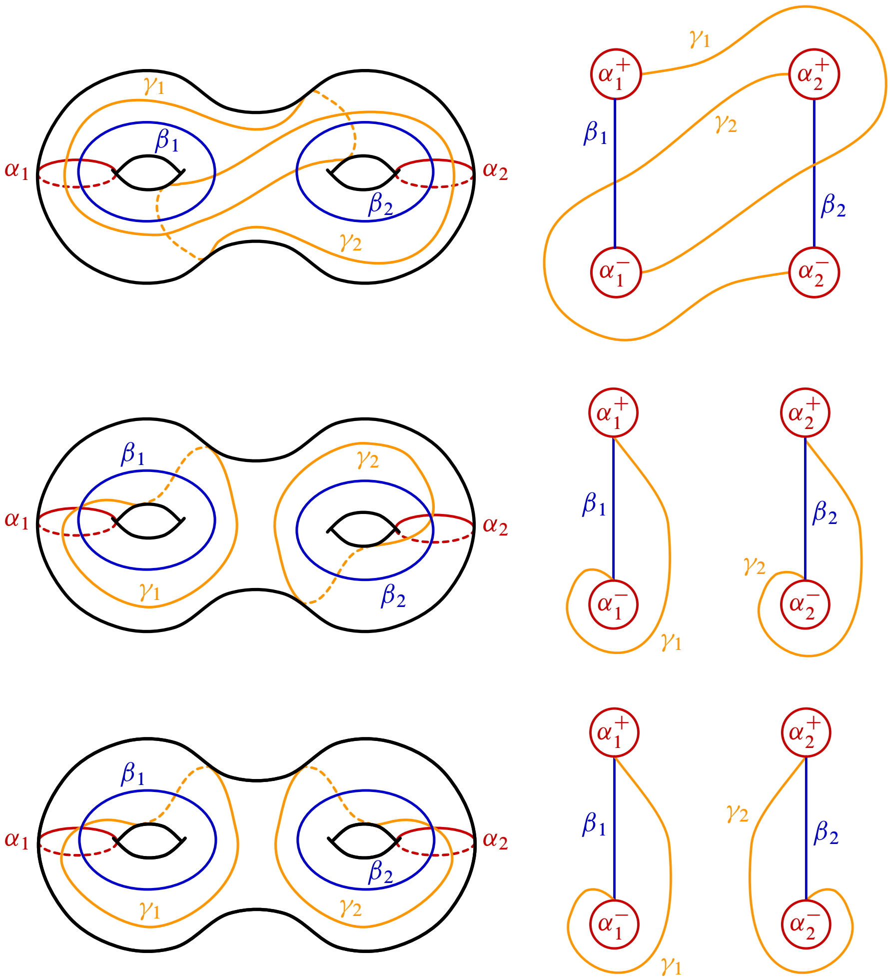

Figure 2: The three standard (2,0)–trisection diagrams: $S^2 \times S^2$ (top), $\mathbb{CP}^2 \# \mathbb{CP}^2$ (middle) and $\mathbb{CP}^2 \# \overline{\mathbb{CP}}^2$ (bottom). By Theorem 1.3, every $(2,0)$–trisection is represented by one of these three diagrams.

Corollary 1.4 *Suppose that L is a tunnel number one link that admits an S^3 surgery with integral slope. Then there is a genus-two Heegaard surface Σ for S^3 containing L such that a series of handle slides on L contained in Σ converts L to a (± 1)–framed unlink or a 0–framed Hopf link.*

Note that the first examples of nontrivial tunnel number one links with nontrivial surgeries were given by Berge [1], and have since been studied extensively by Ishihara [12], Mayrand [16], Matsuda, Ozawa and Shimokawa [15], Ochiai [20] and Teragaito [21].

In Section 7, we discuss how Corollary 1.4 relates to the search for exotic, simply connected 4–manifolds with $b_2 = 2$. We also remark on the similarity between Corollary 1.4 and the Generalized Property R Conjecture in Gompf, Scharlemann and Thompson [7]. (See [17] for more details.)

Organization In Section 2, we introduce notation and terminology which will be used in the rest of the paper. In Section 3, we define a complexity measure for $(2,0)$–trisection diagrams and prove some preliminary results about this complexity. Section 4 includes a digression into a special type of Heegaard diagram for S^3, which we will use to further reduce the complexity of an arbitrary trisection diagram in Section 5. We prove the main theorem about $(2,0)$–trisections in Section 6 and conclude by classifying $(2,1)$–trisections and establishing Corollary 1.4 in Section 7.

Acknowledgements The authors would like to thank Ken Baker and David Gay for their interest in this project and for their insightful comments at the outset, as well as Cameron Gordon for helpful discussions related to the intricacies of the arguments presented below. We would also like to thank the anonymous referee for a careful reading of the manuscript. The first author is supported by NSF grant DMS-1400543, and the second author is partially supported by NSF grant DMS-1203988.

2 Preliminary set-up

Manifolds are orientable and connected unless otherwise specified. For the remainder of the paper, we fix a genus-two surface Σ and let $\nu(\cdot)$ denote an open regular neighborhood in Σ. We use the term *curve* to mean an isotopy class of simple closed curves in Σ. Given two distinct curves c_1 and c_2 in Σ, we assume that c_1 and c_2 have been isotoped so that $|c_1 \cap c_2|$ is minimal, and we let $\iota(c_1, c_2)$ denote this geometric intersection number. By orienting c_1 and c_2, we can also compute the algebraic intersection number, denoted $c_1 \cdot c_2$, by counting the signed intersection number of c_1 and c_2.

A *cut system* α for Σ is a pair of disjoint curves α_1 and α_2 that cut Σ into a four-punctured sphere, which we will denote Σ_α. In a slight abuse of language, we refer to Σ_α as both a four-punctured sphere and a sphere with four boundary components. A *Heegaard diagram* (Σ, α, β) consists of two cut systems α and β for Σ, whereas a *trisection diagram* $(\Sigma, \alpha, \beta, \gamma)$ consists of three cut systems α, β and γ for Σ with the additional restrictions described above. We often suppress Σ and simply write (α, β) and (α, β, γ) for Heegaard and trisection diagrams, respectively. Given a Heegaard diagram (α, β), we may construct a 3–manifold by gluing 3–dimensional

2–handles along α and β on opposite sides of a product neighborhood $\Sigma \times I$ of Σ and capping off the resulting 2–sphere boundary components with 3–balls. Similarly, given a trisection diagram $(\Sigma, \alpha, \beta, \gamma)$, we may construct a 4–manifold by taking a 4–dimensional regular neighborhood $\Sigma \times D$ of Σ and attaching 4–dimensional 2–handles along $\alpha \times \{1\}$, $\beta \times \{e^{2\pi i/3}\}$ and $\gamma \times \{e^{4\pi i/3}\}$, where D is viewed as the complex unit disk and the framings of the handle attachments are given by $\Sigma \times \{p\}$. Finally, we cap off boundary components with 3– and 4–handles, which can be done in a unique way by [14].

We begin by stating several standard lemmata dealing with genus-two Heegaard diagrams of S^3. We say that a Heegaard diagram (α, β) for S^3, where $\alpha = \{\alpha_1, \alpha_2\}$ and $\beta = \{\beta_1, \beta_2\}$, is *trivial* if $\iota(\alpha_i, \beta_j) = \delta_{ij}$. Trivial Heegaard diagrams are unique (up to homeomorphism). Given a Heegaard diagram (α, β), we define the *intersection matrix* $M(\alpha, \beta)$ of (α, β) to be

$$M(\alpha, \beta) = \begin{bmatrix} \alpha_1 \cdot \beta_1 & \alpha_1 \cdot \beta_2 \\ \alpha_2 \cdot \beta_1 & \alpha_2 \cdot \beta_2 \end{bmatrix}.$$

Lemma 2.1 *If (α, β) is a Heegaard diagram for S^3, then $|\det(M(\alpha, \beta))| = 1$.*

It follows immediately that if (α, β) is a Heegaard diagram for S^3 satisfying $\iota(\alpha, \beta) = 2$, then (α, β) is the trivial diagram. Now, suppose that α_1, α_2 and α_3 are pairwise disjoint curves in Σ, and let $\alpha = \{\alpha_1, \alpha_2\}$ and $\alpha' = \{\alpha_2, \alpha_3\}$. We say that α and α' are related by a *handle slide*.

Lemma 2.2 [13] *Two Heegaard diagrams (α, β) and (α', β') determine the same Heegaard splitting if and only if there are sequences of handle slides taking α to α' and β to β'.*

We now turn our attention to trisection diagrams. The first lemma is due to Gay and Kirby.

Lemma 2.3 [6] *Two trisection diagrams (α, β, γ) and $(\alpha', \beta', \gamma')$ determine the same trisection if and only if there are sequences of handle slides taking α to α', β to β', and γ to γ'.*

We say that a $(2,0)$–trisection diagram (α, β, γ) is *standard* if $\iota(\alpha, \beta) = \iota(\alpha, \gamma) = \iota(\beta, \gamma) = 2$. In Figure 2, Σ_α is shown for each of the three standard $(2,0)$–trisection diagrams. The main theorem presented in this paper asserts that for any $(2,0)$–trisection diagram (α, β, γ), there is a sequence of handle slides taking (α, β, γ) to one of these three standard diagrams.

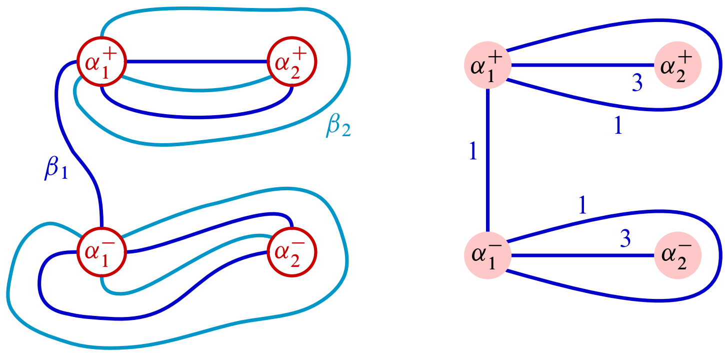

Figure 3: An example of the Whitehead graph $\Sigma_\alpha(\beta)$ (left) and the corresponding reduced Whitehead graph $G_\alpha(\beta)$ (right) obtained from a Heegaard splitting (α, β).

For $\alpha = \{\alpha_1, \alpha_2\}$ and $\mathcal{C}$ a collection of curves in Σ, intersection $\mathcal{C} \cap \Sigma_\alpha$ is a collection of properly embedded essential arcs. Viewing the boundary components of Σ_α as four fat vertices and the essential arcs as edges, we may consider $\mathcal{C} \cap \Sigma_\alpha$ as a (topological) graph, the *Whitehead graph* of $\mathcal{C}$ with respect to α, denoted $\Sigma_\alpha(\mathcal{C})$. By collapsing each boundary component of Σ_α to a vertex and collections of parallel edges in $\Sigma_\alpha(\mathcal{C})$ to a single edge, we construct the *reduced Whitehead graph* $G_\alpha(\mathcal{C})$. To avoid confusion, we will refer to edges in $\Sigma_\alpha(\mathcal{C})$ as *arcs* and edges in $G_\alpha(\mathcal{C})$ as *edges*. If an arc a in $\Sigma_\alpha(\mathcal{C})$ corresponds to an edge e in $G_\alpha(\mathcal{C})$, we say that a is *in the edge* e. The *weight* of an edge $e \in G_\alpha(\mathcal{C})$, denoted $w_{\mathcal{C}}(e)$, is the number of arcs in e. If $\mathcal{C}' \subset \mathcal{C}$, then $\Sigma_\alpha(\mathcal{C}')$ is a subgraph of $\Sigma_\alpha(\mathcal{C})$ and $G_\alpha(\mathcal{C}')$ is a subgraph of $G_\alpha(\mathcal{C})$. In this case we let $w_{\mathcal{C}'}(e)$ denote the weight of e as an edge in $G_\alpha(\mathcal{C}')$. We use $\alpha_i^\pm$ $(i = 1, 2)$ to denote the vertices of both $\Sigma_\alpha(\mathcal{C})$ and $G_\alpha(\mathcal{C})$.

Nonisotopic arcs a and b in Σ_α contained in curves c and d in Σ can intersect in two fundamentally different ways. For an intersection point $p \in a \cap b$, if there is an isotopy of c and d in Σ which pushes p through $\alpha_i^\pm$ without introducing additional intersections of c or d with α, we call p an *inessential point of intersection based in* $\alpha_i^\pm$. Note that an outermost such intersection point is a vertex in a triangle cobounded by subarcs of a, b and a boundary component $\alpha_i^\pm$. A point $p \in a \cap b$ which is not inessential is called an *essential point of intersection*. See Figure 4. Observe that essential points of intersection are present in $G_\alpha(c \cup d)$ but we lose all information about inessential points of intersection when passing from $\Sigma_\alpha(c \cup d)$ to $G_\alpha(c \cup d)$.

An arc a in Σ_α with endpoints in disjoint boundary components is called a *seam*. Otherwise a is called a *wave*. Seams in Σ_α may be parametrized by the extended rational numbers $\overline{\mathbb{Q}}$ in the following way: Choose four pairwise disjoint seams which

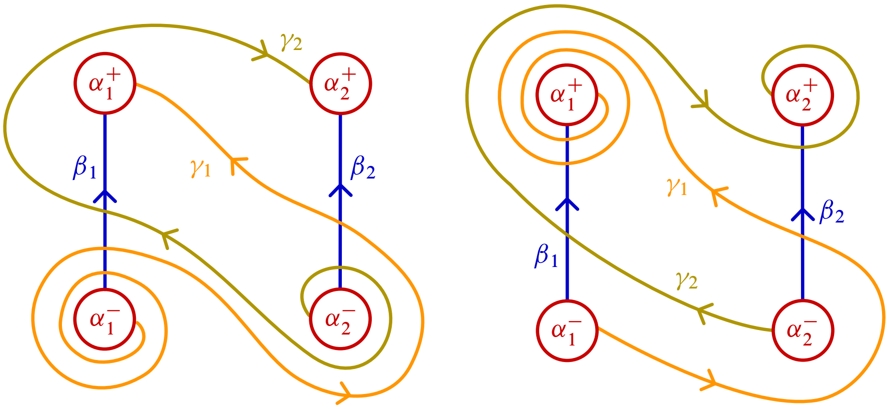

Figure 4: At left, an example of $\Sigma_\alpha(\beta,\gamma)$ such that $\beta\cap\gamma$ contains a total of three inessential intersection points, all based in the vertices α_i^-. At right, curves in γ have been isotoped so that inessential intersections are based in the vertices α_i^+.

cut Σ into two squares with their vertices removed. Assign the left and right sides of the square the slope $\frac{1}{0}$ and the top and bottom the slope $\frac{0}{1}$. Any other seam may be isotoped to have constant rational slope in each of the squares, and we assign this slope to the seam.

We call a seam in Σ_α that connects either α_1^+ or α_1^- to either α_2^+ or α_2^- a *cross-seam*, since it travels across Σ_α from one pair of boundary components to the other. Cross-seams come in pairs: There is a hyperelliptic involution J of Σ that leaves every curve invariant [9]. Thus, the restriction of J to Σ_α is a homeomorphism of Σ_α which maps $\alpha_i^\pm$ to $\alpha_i^\mp$. It follows that if c is a curve in Σ and a_+ is a cross-seam of $c\cap\Sigma_\alpha$ connecting to α_1^+ to $\alpha_2^\pm$, then since $J(c)=c$, we have that $a_-=J(a_+)$ is a cross-seam of $c\cap\Sigma_\alpha$ connecting α_1^- to $\alpha_2^\mp$. See Figure 9.

Next, we show that the slopes of seams determine the number of times they intersect essentially. For each essential arc $a\subset\Sigma_\alpha$, there is a unique essential curve in Σ_α, which we will henceforth call Δ_a, such that $a\cap\Delta_a=\varnothing$. See Figure 5.

Lemma 2.4 *If a and b are seams in Σ_α parametrized by slopes $\frac{w}{x}$ and $\frac{y}{z}$, respectively, and a and b have endpoints on k common boundary components of $\partial\Sigma_\alpha$ (where $k\in\{0,1,2\}$), then a and b intersect essentially in $\frac{1}{2}(|wz-xy|-k)$ points.*

Proof It is well-known that $\iota(\Delta_a,\Delta_b)=2|wz-xy|$. Each common boundary component of a and b contributes two points of intersection to $\Delta_a\cap\Delta_b$, and each

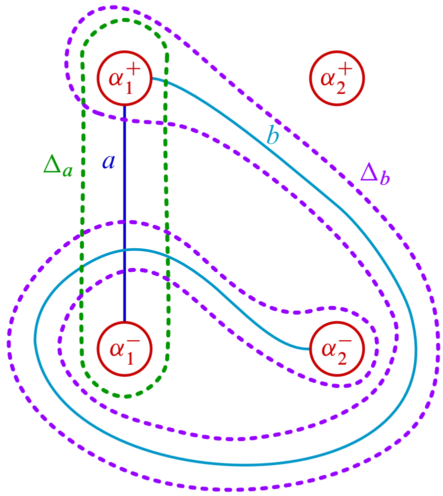

Figure 5: An example depicting the relationship between $\iota(a,b)$ and $\iota(\Delta_a,\Delta_b)$ for arcs a of slope $\frac{1}{0}$ and b of slope $-\frac{1}{3}$, as discussed in the proof of Lemma 2.4

essential point of intersection of a and b contributes four points of intersection to $\Delta_a \cap \Delta_b$. See Figure 5. If there are q essential points of $a \cap b$, it follows that

$$2|wz - xy| = 2k + 4q,$$

which easily yields the desired statement. □

Given a Heegaard diagram (α,β), an α*–wave* for (α,β) is wave ω properly embedded in Σ_α such that $\omega \cap \beta = \varnothing$. A β*–wave* is defined similarly. If such an arc exists, we say that (α,β) contains a wave *based in* α (resp. β) and that $\Sigma_\alpha(\beta)$ (resp. $\Sigma_\beta(\alpha)$) contains a wave. Let $\alpha = \{\alpha_1,\alpha_2\}$ and $\alpha_3 = \Delta_\omega$, and suppose that $\partial\omega \subset \alpha_1^\pm$. Then $\alpha' = \{\alpha_2,\alpha_3\}$ is related to α by a handle slide, and we say that the diagram (α',β) is the result of *wave surgery* (or just *surgery*, for short) on α along ω.

Proposition 2.5 [11] *Let (α,β) be a Heegaard diagram for S^3 that contains a wave based in α, and let (α',β) be the diagram obtained by wave surgery on α. Then $\iota(\alpha',\beta) < \iota(\alpha,\beta)$.*

Thus, it is desirable for a nontrivial Heegaard diagram to contain a wave, as it allows for a straightforward simplification of the diagram. Unfortunately, this is not the case for Heegaard diagrams of S^3 of arbitrary genera [19]. However, the situation in genus two is much more manageable, as is shown by the following theorem of Homma, Ochiai and Takahashi.

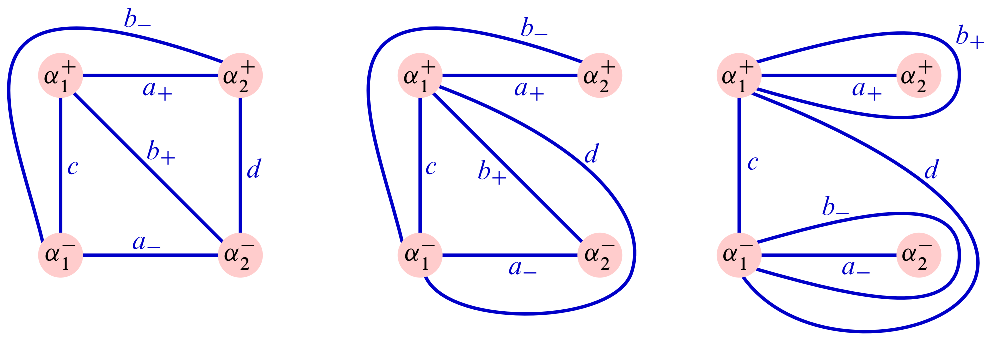

Figure 6: The three possible reduced Whitehead graphs $G_\alpha(\beta)$ corresponding to a genus-two Heegaard splitting: from left to right, Types I, II and III

Theorem 2.6 [11] *Every nontrivial genus-two Heegaard diagram for S^3 contains a wave.*

Wave surgery has a dual notion, known as *banding*. Let $\alpha' = \{\alpha_2, \alpha_3\}$ and suppose υ is an arc with one endpoint on each of α_2 and α_3 and whose interior is disjoint from α' (we call υ a *band*). Then one component α_1 of $\partial\overline{\nu(\alpha' \cup \upsilon)}$ is not isotopic to α_2 or α_3, and we that say α_1 is the result of banding α_2 to α_3. For every wave ω such that wave surgery on $\alpha = \{\alpha_1, \alpha_2\}$ yields $\alpha' = \{\alpha_2, \alpha_3\}$, there is a dual band υ such that α_1 is the result of banding α_2 to α_3; thus, banding and wave surgery are inverse processes.

We now state several general facts about the reduced Whitehead graphs corresponding to genus-two Heegaard diagrams. First, every such graph is isomorphic to one of the graphs shown in Figure 6 (see [19]), which we henceforth refer to as Type I, Type II and Type III as labeled in the figure. We will automatically assume that a graph of one of these types inherits the edge labelings shown. Note that if all edge weights are nonzero, then the three families are mutually exclusive; however, if some edge weights are zero there may be graphs that fall into more than one family. In addition, the existence of the involution J implies that $w_\beta(a_+) = w_\beta(a_-)$ and $w_\beta(b_+) = w_\beta(b_-)$ and so we may denote these weights $w_\beta(a)$ and $w_\beta(b)$ without ambiguity.

If a reduced Whitehead graph $G_\alpha(\beta)$ contains exactly four edges and one 4–cycle, we will call $G_\alpha(\beta)$ a *square graph*. Equivalently, a square graph is a Type I graph such that $w_\beta(b) = 0$ and all other edges have nonzero weights. If $G_\alpha(\beta)$ is a graph of any type such that $w_\beta(a), w_\beta(c) \neq 0$ while $w_\beta(b) = w_\beta(d) = 0$, we say that $G_\alpha(\beta)$ is a *c-graph* (as its shape looks like the letter "c").

The next lemma provides basic information about some of the weights in $G_\alpha(\beta)$ when (α, β) is a diagram for S^3.

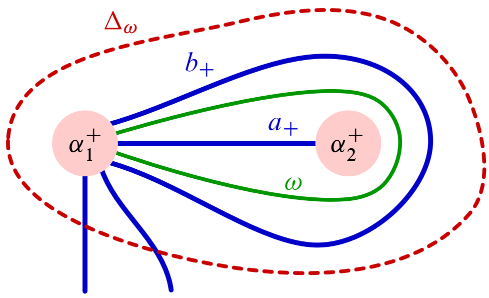

Figure 7: An α–wave ω in a Type III graph. If the weight of a or b is nonzero, then $\iota(\Delta_\omega, \beta) < \iota(\alpha_1, \beta)$.

Lemma 2.7 [11, Lemma 9] *If (α, β) is a nontrivial diagram for S^3, then $\Sigma_\alpha(\beta)$ contains a cross-seam.*

Observe that if (α, β) contains a wave based in α, then the graph $G_\alpha(\beta)$ must be of Type III. A graph of Type I or II that admits a wave must also fall under the Type III category. Figure 7 shows the obvious wave in this case. Equivalently, (α, β) contains a wave based in α if and only if a vertex in $G_\alpha(\beta)$ has valence one (recall that $G_\alpha(\beta)$ is a reduced graph). Note that Theorem 2.6 implies that if (α, β) is a Heegaard diagram for S^3, then there is a wave based in *either* α or β. However, if $G_\alpha(\beta)$ is a c-graph, we can say something stronger:

Lemma 2.8 *If (α, β) is a diagram for S^3 and $G_\alpha(\beta)$ is a c-graph, then (α, β) contains a wave based in α and a wave based in β.*

Proof In this case, $\Sigma \setminus \nu(\alpha \cup \beta)$ contains some number of rectangular regions and one 12–sided region, which we will call U, where the boundary of U alternates between a set of six arcs in α and a set of six arcs in β. See Figure 8. A choice of orientation α and β orients each arc in ∂U clockwise or counterclockwise. Moreover, each α–arc comes from either α_1 or α_2. Therefore, there are four labeling options for each α–arc. However, since there are six α–arcs, at least two have a common labeling, and an arc ω_α in U connecting two such arcs is a wave based in α. See Figure 8 (right). A parallel argument shows that there is a wave ω_β in U based in β. □

In the case-by-case arguments that follow, we attempt to rule out the existence of certain waves for a Heegaard diagram, so we lay the groundwork for that strategy here. If two oriented curves meet in two points, we say that these points are *coherently oriented* if they have the same sign and *oppositely oriented* if they have opposite sign. Note that this definition is independent of the choice of orientations on the curves. If (α, β) is a

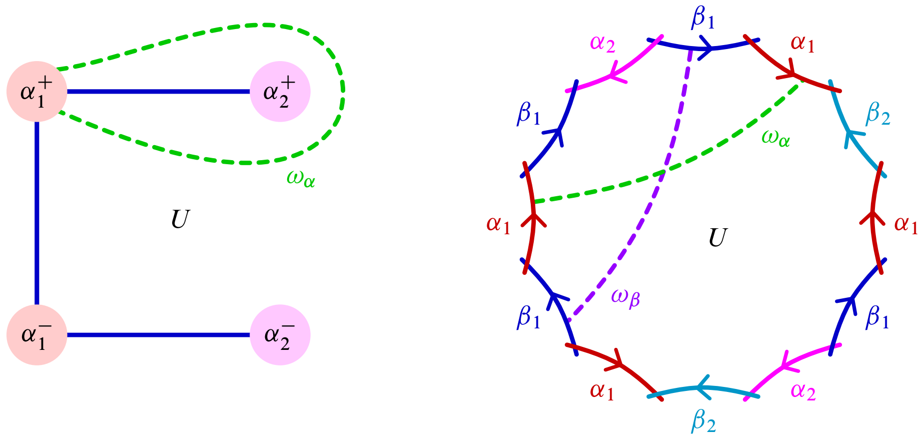

Figure 8: When $G_\alpha(\beta)$ is a c-graph, there are waves in α and β. The obvious wave ω_α is shown (left), and the rendering of the 12–gon U (right) implies the existence of ω_β.

Heegaard diagram and there is a choice of orientations on curves in α and β so that all points of $\alpha \cap \beta$ are coherently oriented, we call (α, β) a *positive Heegaard diagram*.

Suppose that (α, β) is a Heegaard diagram and that β_1 contains an arc a that meets α_1 only in its oppositely oriented endpoints and meets α_2 in one point in its interior. We call an arc a satisfying these conditions a *wave-busting arc*. Similarly, if β contains two arcs c and d such that c meets α_1 only in its coherently oriented endpoints and avoids α_2, while d meets α_2 only in its coherently oriented endpoints and avoids α_1, we say that the pair (c, d) is a *wave-busting pair*. The next lemma makes clear our use of the term "wave-busting".

Lemma 2.9 *If (α, β) is a nontrivial Heegaard diagram such that β contains a wave-busting arc a or a wave-busting pair (c, d), then (α, β) does not contain a wave based in α.*

Proof A wave-busting arc a contributes two cross-seams to $G_\alpha(\beta)$, one from $\alpha_1^\pm$ to α_2^+ and one from $\alpha_1^\pm$ to α_2^-. It follows by the symmetry of $G_\alpha(\beta)$ that each vertex has valence at least two, and thus $G_\alpha(\beta)$ cannot be a graph of Type III. On the other hand, a wave-busting pair contributes seams connecting α_1^+ to α_1^- and α_2^+ to α_2^-, again implying that $G_\alpha(\beta)$ is not a Type III graph. □

To conclude this section, we discuss how homeomorphisms of Σ_α can be used to reduce the number of cases in the arguments that follow. The five homeomorphisms

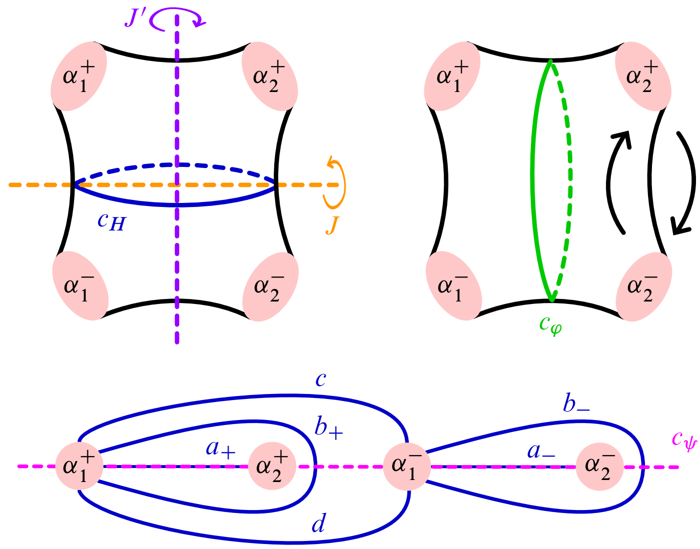

Figure 9: The five homeomorphisms of the four-punctured sphere Σ_α that will be useful throughout the paper. First, J and J' are rotations through π radians about the axes shown (top left), and H is a reflection across the plane that intersects the sphere in c_H. The homeomorphism φ is a half-twist fixing c_φ and exchanging α_2^+ and α_2^- (top right), and the action of ψ is shown as a reflection across the plane intersecting the page in the line c_ψ (bottom).

most relevant to our proof are shown in Figure 9. The mapping class group of a four-punctured sphere is generated by $\mathrm{PGL}_2(\mathbb{Z})$ and two involutions [2], the first of which is the hyperelliptic involution J discussed above, and the second of which we will call J' and is depicted in Figure 9 (top left). Note that J and J' preserve the slopes of arcs in Σ_α, so to determine the behavior of slopes of arcs with respect to a given homeomorphism σ, we need only specify an element of $\mathrm{PGL}_2(\mathbb{Z})$, which we will call M_σ.

Let $c_\varphi = \Delta_a$ for a seam a of slope $\frac{1}{0}$, and let the homeomorphism φ be given by fixing $\alpha_1^\pm$ and c_φ and performing a half-twist on $\alpha_2^\pm$ which exchanges these two punctures, as in Figure 9 (top right). The matrices for φ and φ^{-1} are given by

$$M_\varphi = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad M_{\varphi^{-1}} = \begin{bmatrix} 1 & -1 \\ 0 & 1 \end{bmatrix}.$$

The remaining two homeomorphisms reverse the orientation of Σ_α. Let H be the horizontal reflection fixing the curve c_H shown in Figure 9 (top left), and let ψ be a reflection of Σ_α fixing the four arcs denoted c_ψ in Figure 9 (bottom). Note that ψ preserves a Type III graph; edges $a_\pm$ and $b_\pm$ are fixed, while edges c and d are exchanged. The matrices for H and ψ are given by

$$M_H = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \quad \text{and} \quad M_\psi = \begin{bmatrix} -1 & 0 \\ 2 & 1 \end{bmatrix}.$$

3 The complexity of a $(2,0)$–trisection diagram

To classify $(2,0)$–trisections, we will show that a trisection diagram which minimizes a prescribed complexity over all diagrams corresponding to a fixed $(2,0)$–trisection must be a standard diagram. Define the *complexity* $\mathcal{C}(\alpha,\beta,\gamma)$ of a trisection diagram (α,β,γ) to be the triple given by

$$\mathcal{C}(\alpha,\beta,\gamma) = (\iota(\alpha_2,\gamma), \iota(\alpha_1,\gamma), \iota(\beta,\gamma)).$$

For a $(2,0)$–trisection $X = X_1 \cup X_2 \cup X_3$, we call a trisection diagram *minimal* if it minimizes the complexity $\mathcal{C}$ (with the dictionary ordering) among diagrams satisfying $\iota(\alpha,\beta) = 2$. A plurality of this paper (Sections 3–5) is devoted to proving the following proposition, which we will use later to classify $(2,0)$–trisections in Section 6.

Proposition 3.1 *If (α,β,γ) is minimal, then $\iota(\alpha,\gamma) = 2$.*

To prove the proposition, we will suppose by way of contradiction that (α,β,γ) is minimal, but (α,γ) is not trivial, so that $\iota(\alpha,\gamma) > 2$. We may also suppose that (β,γ) is not trivial; otherwise $\mathcal{C}(\beta,\gamma,\alpha) < \mathcal{C}(\alpha,\beta,\gamma)$ and (α,β,γ) is not minimal. One case is quite easy to rule out.

Lemma 3.2 *If $G_\gamma(\alpha)$ is a Type III graph, then (α,β,γ) is not minimal.*

Proof In this case, $\Sigma_\gamma(\alpha)$ contains a wave, and we let γ' be the result of surgery on γ along this wave, so that γ and γ' are related by a handle slide. With Figure 7 in mind, we compute

$$\iota(\alpha_2,\gamma) = 2w_{\alpha_2}(a) + 2w_{\alpha_2}(b) + w_{\alpha_2}(c) + w_{\alpha_2}(d) \tag{1}$$

$$\iota(\alpha_2,\gamma') = w_{\alpha_2}(a) + w_{\alpha_2}(c) + w_{\alpha_2}(d). \tag{2}$$

Thus, if $w_{\alpha_2}(a)$ or $w_{\alpha_2}(b)$ is nonzero in $G_\gamma(\alpha)$, then $\iota(\alpha_2,\gamma') < \iota(\alpha_2,\gamma)$ and $\mathcal{C}(\alpha,\beta,\gamma') < \mathcal{C}(\alpha,\beta,\gamma)$. Otherwise, we have $\iota(\alpha_2,\gamma) = \iota(\alpha_2,\gamma')$. However, since γ' is surgery on γ along a wave, $\iota(\alpha,\gamma') < \iota(\alpha,\gamma)$, which implies that $\iota(\alpha_1,\gamma') < \iota(\alpha_1,\gamma)$ and, once again, $\mathcal{C}(\alpha,\beta,\gamma') < \mathcal{C}(\alpha,\beta,\gamma)$. In either case, (α,β,γ) is not minimal. □

Using equations (1) and (2), we also have the next lemma.

Lemma 3.3 *If $G_\gamma(\alpha_2)$ is a Type III graph that contains a cross-seam, then (α,β,γ) is not minimal.*

Proof In this case, we have $w_{\alpha_2}(a)\neq 0$, and so (α,β,γ) is not minimal by the proof of Lemma 3.2. □

By Theorem 2.6 and the assumption that (α,γ) is not trivial, we have that (α,γ) contains a wave ω. If ω is based in γ, then $G_\gamma(\alpha)$ is a Type III graph, so (α,β,γ) is not minimal by Lemma 3.2. Thus, we may assume that ω is based in α, and $G_\alpha(\gamma)$ is a Type III graph. Since (α,β,γ) is minimal, we have that $\iota(\alpha_2,\gamma)\leq\iota(\alpha_1,\gamma)$; otherwise we could change the labeling of $\alpha=\{\alpha_1,\alpha_2\}$ to reduce complexity. Therefore, we can assume that $\partial\omega\subset\alpha_1$.

We parametrize arcs in $G_\alpha(\gamma)$ by assigning the seams connecting α_1^+ to α_1^- the slopes $\frac{1}{0}$ and $-\frac{1}{2}$, and the seams connecting $\alpha_1^\pm$ to $\alpha_2^\pm$ the slope $\frac{0}{1}$. Since $\iota(\alpha,\beta)=2$, each intersection $\beta_i\cap\Sigma_\alpha$ is a single seam connecting α_i^+ and α_i^- for $i=1,2$, and these two seams have the same slope. Let $\frac{m}{n}$ denote the slope of these seams, noting that m must be odd and n must be even, and we suppose that $m>0$. We also assume that $\iota(\alpha_i,\beta_j)=\delta_{ij}$. In the case that $m=1$, there is a curve β_3 such that $\beta_3\cap\Sigma_\alpha$ consists of two cross-seams of slope $\frac{0}{1}$ and $\beta_3\cap\beta=\varnothing$. See Figure 10. We call the curve β_3 the 0*–replacement* for β.

Lemma 3.4 *If $m=1$, then (α,β,γ) is not minimal.*

Proof Let $\alpha'=\{\alpha_3,\alpha_2\}$ be the result of surgery on α along the wave ω, let β_3 be the 0–replacement for β, and let $\beta'=\{\beta_1,\beta_3\}$. Then α and β are related to α' and β' by handle slides, and $\iota(\alpha',\beta')=2$. However, $\iota(\alpha_3,\gamma)<\iota(\alpha_1,\gamma)$ and thus $\mathcal{C}(\alpha',\beta',\gamma)<\mathcal{C}(\alpha,\beta,\gamma)$. See Figure 10. □

We suppose for the remainder of this section that $m\geq 3$. Our plan of attack is as follows: We show that, in most cases, (β,γ) contains a wave based in γ and that (α,β,γ) is not minimal by Lemma 3.6. In the exceptional case, we will prove in Section 5 that (β,γ) is not a Heegaard diagram for S^3, a contradiction.

Before we proceed, we require a technical lemma, which we state in greatest generality for use in later sections.

Lemma 3.5 *Suppose that (δ,ϵ) is a Heegaard diagram such that $\iota(\delta_1,\epsilon_2)=0$ and $G_\delta(\epsilon)$ is a square graph. Then (δ,ϵ) is not a Heegaard diagram for S^3.*

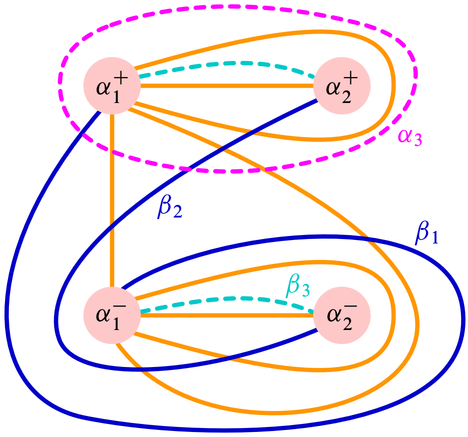

Figure 10: The case in which $m = 1$, so that arcs of β_i have slope $\frac{1}{n}$ (shown with $n = 2$). Replacing α_1 with α_3 and β_2 with the 0–replacement β_3 lowers the complexity of the diagram.

Proof We label the edges of the Type I graph $G_\delta(\epsilon)$ as in Figure 6, noting that ϵ_2 contributes to only one edge, d, in $G_\delta(\epsilon)$. Let $k_1 = w_{\epsilon_2}(d)$. Since $G_\delta(\epsilon)$ is a square graph, we may orient δ and ϵ so that all intersection points have the same sign; thus, (δ, ϵ) is a positive Heegaard diagram. It follows that $\iota(\delta_i, \epsilon_j) = \delta_i \cdot \epsilon_j$. Since all edges of $G_\delta(\epsilon)$ have nonzero weights, ϵ_1 intersects δ_2 in $k_2 = w_\epsilon(a) > 0$ points and δ_1 in $k_3 = w_\epsilon(a) + w_\epsilon(c) > 1$ points. This implies that the intersection matrix $M(\delta, \epsilon)$ is

$$M(\delta, \epsilon) = \begin{bmatrix} k_3 & 0 \\ k_2 & k_1 \end{bmatrix},$$

where $|\det(M(\delta, \epsilon))| = k_1 k_3 > 1$. Thus, (δ, ϵ) is not a Heegaard diagram for S^3 by Lemma 2.1. □

Now, we return to (α, β, γ). If β^* is a seam in $\Sigma_\alpha(\beta)$ connecting α_i^+ to α_i^-, then the slope of β^* is $\frac{m}{n}$. A seam γ^* in $\Sigma_\alpha(\gamma)$ connecting α_1^+ to α_1- has slope $\frac{1}{0}$ or $-\frac{1}{2}$, while a cross-seam has slope $\frac{0}{1}$. It follows from Lemma 2.4 that β^* intersects γ^* essentially. The same is true if γ^* is a wave arc. Recall that we have assumed that (β, γ) is not trivial, so by Theorem 2.6, (β, γ) contains a wave.

Lemma 3.6 *If (β, γ) contains a wave based in γ, then (α, β, γ) is not minimal.*

Proof If (β, γ) contains a wave based in γ, then $G_\gamma(\beta)$ is a Type III graph, and thus each vertex $\gamma_2^\pm$ is incident to only one edge $a_\pm$ in $G_\gamma(\beta)$. Isotope α so that any inessential intersections of α and β in Σ_γ based in $\gamma_2^\pm$ (there are at most two of them)

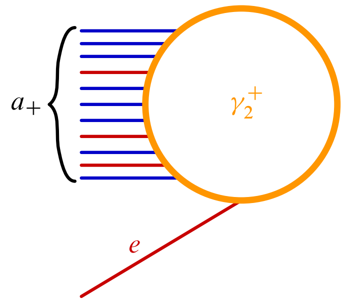

Figure 11: A local picture of Σ_γ near γ_2^+. All arcs of $\Sigma_\gamma(\beta)$ and all but one arc of $\Sigma_\gamma(\alpha)$ meeting γ_2^+ are in the edge a_+, and the one exceptional arc of α is in the edge e.

are based in γ_2^-. We note that γ_2^+ does not contain two consecutive intersections with α. Otherwise, the subarc connecting such intersections corresponds to an arc γ^* in $\Sigma_\alpha(\gamma)$ which avoids β, contradicting that every such arc γ^* intersects an arc of β essentially in Σ_α.

It follows that in $\Sigma_\gamma(\alpha,\beta)$, all α–arcs meeting γ_2^+ with possibly one exception are situated between β–arcs in the edge a_+. In other words, in the reduced graph $G_\gamma(\alpha)$, the vertex γ_2^+ is incident to at most two edges, one of which has the same slope as a_+ in $G_\gamma(\beta)$, and so we will call this edge a_+, and the other of which we will call e. See Figure 11. Note that $w_\alpha(e)=1$. We split the remainder of the argument into four cases, based on possible values of $w_{\alpha_2}(a)$ and $w_{\alpha_2}(e)$.

Case 1 $w_{\alpha_2}(a)\neq 0$ and $w_{\alpha_2}(e)=0$. In this case, the vertex γ_2^+ has valence one in $G_\gamma(\alpha_2)$ and $G_\gamma(\alpha_2)$ contains the cross-seam a_+. By Lemma 3.3, the diagram (α,β,γ) is not minimal.

Case 2 $w_{\alpha_2}(a)=0$ and $w_{\alpha_2}(e)\neq 0$. If e is a cross-seam, then by the same argument as in Case 1, (α,β,γ) is not minimal. Otherwise, e connects γ_2^+ to γ_2^-. By the argument above, $\gamma_2^\pm$ is incident to exactly two edges $a^\pm$ and e in $G_\gamma(\alpha)$. If $\gamma_1^\pm$ has valence one in $G_\gamma(\alpha)$, then $G_\gamma(\alpha)$ is a c-graph and there is a wave for (α,γ) based in γ by Lemma 2.8 and thus (α,β,γ) is not minimal by Lemma 3.2. If γ_1^+ has valence greater than one, there is another edge c incident to γ_1^+ in $G_\gamma(\alpha)$, and c must connect γ_1^+ to γ_1^-. It follows that $G_\gamma(\alpha)$ is a Type I graph with only one cross-seam; equivalently, $G_\gamma(\alpha)$ is a square graph, and the assumption $w_{\alpha_2}(a)=0$ implies that $\iota(\alpha_2,\gamma_1)=0$. However, in this case Lemma 3.5 implies that (α,γ) is not a Heegaard diagram for S^3, a contradiction.

Case 3 $w_{\alpha_2}(a)\neq 0$ and $w_{\alpha_2}(e)\neq 0$. Here, $w_{\alpha_2}(e)=1$ and thus $w_{\alpha_1}(e)=0$. Let $\gamma_3=\Delta_a$ and $\gamma'=\{\gamma_2,\gamma_3\}$ so that γ is related to γ' by a handle slide. There are two

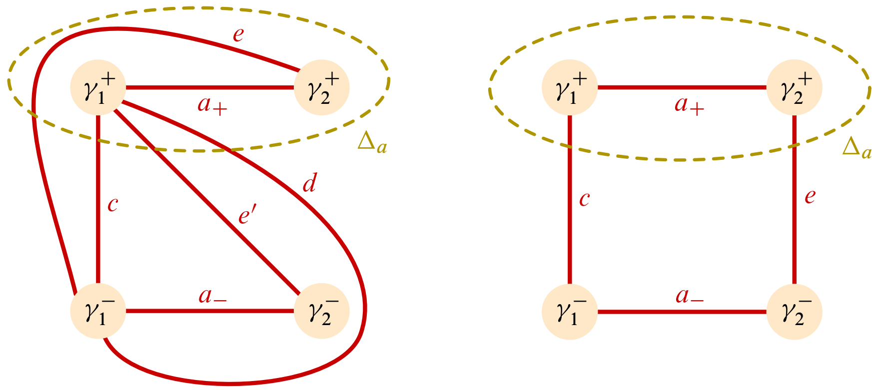

Figure 12: The two subcases of Case 3 considered in the proof of Lemma 3.6: $G_\gamma(\alpha_2)$ is a Type II graph (left), and $G_\gamma(\alpha_2)$ is a square graph (right). In each case, replacing γ_1 by $\gamma_3 = \Delta_a$ reduces complexity.

subcases to consider; refer to Figure 12. Suppose first that e is a cross-seam, so that there is another cross-seam e' which meets γ_2^- by the symmetry of $G_\gamma(\alpha)$. Since the valence of γ_2^+ in $G_\gamma(\alpha)$ is two, we have $G_\gamma(\alpha)$ is a Type II graph. Let c and d denote the two distinct edges connecting γ_1^+ to γ_1^-. We compute

$$\iota(\alpha_2,\gamma_1) = w_{\alpha_2}(a) + w_{\alpha_2}(c) + w_{\alpha_2}(d) + 1,$$
$$\iota(\alpha_2,\gamma_3) = w_{\alpha_2}(c) + w_{\alpha_2}(d) + 2.$$

It follows from $w_{\alpha_2}(a) \neq 0$ that $\iota(\alpha_2,\gamma') \leq \iota(\alpha_2,\gamma)$. If this inequality is strict, then (α,β,γ) is not minimal. Otherwise, we assume $\iota(\alpha_2,\gamma') = \iota(\alpha_2,\gamma)$. Since $w_{\alpha_1}(e) = 0$, $G_\gamma(\alpha_1)$ is a Type III graph. Thus, if $w_{\alpha_1}(a) \neq 0$, then $\iota(\alpha_1,\gamma') < \iota(\alpha_1,\gamma)$ so that $\mathcal{C}(\alpha,\beta,\gamma') < \mathcal{C}(\alpha,\beta,\gamma)$. Otherwise, $w_{\alpha_1}(a) = 0$ and thus $\iota(\alpha_1,\gamma') = \iota(\alpha_1,\gamma)$. Now, since γ' is the result of surgery on a wave for (β,γ), we have $\iota(\beta,\gamma') < \iota(\beta,\gamma)$ in addition to $\iota(\alpha_2,\gamma') = \iota(\alpha_2,\gamma)$ and $\iota(\alpha_1,\gamma') = \iota(\alpha_1,\gamma)$; hence $\mathcal{C}(\alpha,\beta,\gamma') < \mathcal{C}(\alpha,\beta,\gamma)$.

In the second case, we suppose that e connects γ_2^+ to γ_2^-. If the valence of $\gamma_1^\pm$ in $G_\gamma(\alpha_2)$ is one, then (α,β,γ) is not minimal by Lemma 3.3. Otherwise, the valence of $\gamma_1^\pm$ is greater than one in $G_\gamma(\alpha)$, and since the valence of γ_2^+ is two, $G_\gamma(\alpha)$ is a Type I graph with one cross-seam; equivalently, $G_\gamma(\alpha)$ is a square graph. Here we let c denote the single edge connecting γ_1^+ and γ_1^-. As above, we compute

$$\iota(\alpha_2,\gamma_1) = w_{\alpha_2}(a) + w_{\alpha_2}(c),$$
$$\iota(\alpha_2,\gamma_3) = w_{\alpha_2}(c) + 1.$$

If $w_{\alpha_2}(a) > 1$, then $\iota(\alpha_2, \gamma') < \iota(\alpha_2, \gamma)$ and (α, β, γ) is not minimal. On the other hand, if $w_{\alpha_2}(a) = 1$, then $\iota(\alpha_2, \gamma) = \iota(\alpha_2, \gamma')$, while $G_\gamma(\alpha_1)$ is a Type III graph, which implies $\mathcal{C}(\alpha, \beta, \gamma') < \mathcal{C}(\alpha, \beta, \gamma)$ as above.

Case 4 $w_{\alpha_2}(a) = w_{\alpha_2}(e) = 0$. If $w_\alpha(a) = w_\alpha(e) = 0$, then Lemma 2.7 implies that (α, γ) is trivial. If one of $w_\alpha(a)$ or $w_\alpha(e)$ is zero, then $G_\gamma(\alpha)$ is a Type III graph and (α, β, γ) is not minimal by Lemma 3.2. Otherwise, $w_{\alpha_1}(e) = 1$ and $w_{\alpha_1}(a) \neq 0$. Suppose first that e is a cross-seam, so that by the symmetry of $G_\gamma(\alpha_1)$, there is another weight-one cross-seam e' which meets γ_2^-. Since the valence of $\gamma_2^\pm$ is two and $G_\gamma(\alpha)$ contains two distinct cross-seams, the graph $G_\gamma(\alpha)$ is a Type II graph (as in Figure 12 (left)). Let c and d denote the edges connecting γ_1^+ to γ_1^-, let $\gamma_3 = \Delta_a$, and let $\gamma' = \{\gamma_2, \gamma_3\}$ so that γ' is related to γ by a handle slide. Then $\iota(\alpha_2, \gamma_3) = \iota(\alpha_2, \gamma_1) = w_{\alpha_2}(c) + w_{\alpha_2}(d)$, and, using $w_{\alpha_1}(e) = w_{\alpha_1}(e') = 1$, we compute

$$\iota(\alpha_1, \gamma_1) = w_{\alpha_1}(a) + w_{\alpha_1}(c) + w_{\alpha_1}(d) + 1,$$
$$\iota(\alpha_1, \gamma_3) = w_{\alpha_1}(c) + w_{\alpha_1}(d) + 2.$$

Thus, if $w_{\alpha_1}(a) > 1$, then $\iota(\alpha_1, \gamma') < \iota(\alpha_1, \gamma)$ so that $\mathcal{C}(\alpha, \beta, \gamma') < \mathcal{C}(\alpha, \beta, \gamma)$. Otherwise, $w_{\alpha_1}(a) = 1$ and thus $\iota(\alpha_1, \gamma') = \iota(\alpha_1, \gamma)$. As above, since γ' is the result of surgery on a wave for (β, γ), we have $\iota(\beta, \gamma') < \iota(\beta, \gamma)$; hence $\mathcal{C}(\alpha, \beta, \gamma') < \mathcal{C}(\alpha, \beta, \gamma)$.

In the second subcase, suppose that e is a seam connecting γ_2^+ to γ_2^-. Since an α_2–arc contributes an edge, call it c, connecting γ_1^+ to γ_1^- in $G_\gamma(\alpha)$, it follows that $G_\gamma(\alpha)$ is a Type I graph with only one cross-seam (see Figure 12 (right)); that is, $G_\gamma(\alpha)$ is a square graph. Moreover, $w_{\alpha_2}(a) = w_{\alpha_2}(e) = 0$ implies $\iota(\alpha_2, \gamma_2) = 0$, so by Lemma 3.5, we have (α, γ) is not a Heegaard diagram for S^3, a contradiction. □

In the two lemmas that follow, we return our attention to $G_\alpha(\beta, \gamma)$. These two lemmas, together with our previous work in this section, prove Proposition 3.1 in all but one exceptional case. This case is significantly more complicated and requires the arguments in Sections 4 and 5 to settle.

Lemma 3.7 *If $w_\gamma(b) \neq 0$ in $G_\alpha(\gamma)$, then (β, γ) does not contain a wave based in β.*

Proof Observe that there is an arc β^* of β_2 which connects α_2^+ to a point p contained in an arc γ_0 in the edge b_+ of $G_\alpha(\gamma)$. See Figure 13 (left). Since $m \geq 3$, we have that β_1 intersects arcs in the edge a essentially, which implies that β_1 meets γ_0 in a nonempty set of pairs of oppositely oriented points. It follows that γ_0 contains a wave-busting arc γ^* containing the point p; hence $G_\beta(\gamma)$ does not contain a wave by Lemma 2.9. □

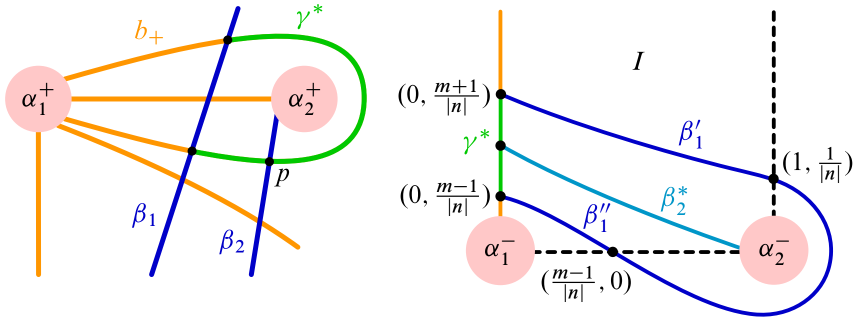

Figure 13: Wave-busting arcs in γ^* in the case that $w_\gamma(b) \neq 0$ (left) and that $w_\gamma(b) = 0$ (right), but the slope $\frac{m}{n} \neq -(n \pm 1)/n$

We now assume that $w_\gamma(b) = 0$. Recall that the homeomorphism ψ described in Section 2 maps $G_\alpha(\gamma)$ onto itself.

Lemma 3.8 *If $\frac{m}{n} \neq -(n \pm 1)/n$, then (β, γ) does not contain a wave based in β or (α, β, γ) is not minimal.*

Proof First, we note that $w_\gamma(a)$, $w_\gamma(c)$ and $w_\gamma(d)$ are not zero; otherwise, $G_\gamma(\alpha)$ is a c-graph and contains a wave by Lemma 2.8, so (α, β, γ) is not minimal by Lemma 3.2. To prove the lemma, we will view Σ_α as two unit squares I and I' glued together along their boundaries with a neighborhood of their vertices removed. We describe points in I with Cartesian coordinates, where $\alpha_1^- = (0,0)$. First, suppose that $\frac{m}{n} < 0$ and $1 < m < |n| - 1$. Let β_2^* denote the arc of $\beta_2 \cap I$ which meets the vertex of I at the coordinates $(1,0)$. Since the slope of β_2^* is $-\frac{m}{|n|}$, the coordinates of the other endpoint of β_2^* are $\left(0, \frac{m}{|n|}\right)$. In I, there are two arcs of β_1, call them β_1' and β_1'', which are adjacent to β_2^*. The endpoints are given by

$$\left(0, \frac{m+1}{|n|}\right) \text{ and } \left(1, \frac{1}{|n|}\right) \text{ for } \beta_1', \quad \left(0, \frac{m-1}{|n|}\right) \text{ and } \left(\frac{m-1}{|n|}, 0\right) \text{ for } \beta_1''.$$

See Figure 13 (right). The assumption that $1 < m < |n| - 1$ ensures that none of these four endpoints is a vertex of I. In addition, we note that in I', the endpoints $\left(1, \frac{1}{|n|}\right)$ and $\left((m-1)/|n|, 0\right)$ are connected by an arc in β_1. Thus, an arc in $\Sigma_\alpha(\gamma)$ which has slope $\frac{1}{0}$ contains a subarc γ^* connecting oppositely oriented points of intersection with β_1, namely $\left(0, (m+1)/|n|\right)$ and $\left(0, (m-1)/|n|\right)$. Further, γ^* meets β_2 precisely once in the point $\left(0, \frac{m}{|n|}\right)$. Therefore, γ^* is a wave-busting arc and (β, γ) does not contain a wave by Lemma 2.9.

For the other cases, we will show that by applying the homeomorphism ψ to Σ_α, we can reduce the argument to the one above. Since ψ preserves $G_\alpha(\gamma)$, we need only determine the action of ψ on the slope of $\frac{m}{n}$ of arcs in β. Suppose that either $\frac{m}{n} > 0$ or $\frac{m}{n} < 0$ and $m > |n| + 1$. We have

$$\psi \begin{bmatrix} m \\ n \end{bmatrix} = \begin{bmatrix} -m \\ 2m+n \end{bmatrix}.$$

If $\frac{m}{n} > 0$, then certainly $-m/(2m+n) < 0$ and $m < 2m+n-1$. Hence, the above argument shows that there exists a wave-busting arc $\gamma^* \subset \gamma$, so (β, γ) does not contain a wave based in γ. Now, suppose that $n < 0$ and $m > |n| + 1$ and let $k = 2m + n = 2m - |n|$, so that $k > m + 1$ and thus $m < k - 1$. Once again, we have that $-m/(2m+n) = -\frac{m}{k}$ is of the form handled by the above argument, and we conclude that (β, γ) does not contain a wave based in β. □

The remaining case to consider in the proof of Proposition 3.1 occurs when $\frac{m}{n} = -(n \pm 1)/n$, and we postpone our examination of this case until Section 5.

4 A brief digression into square graphs

Recall some of the assumptions and conclusions we have made up to this point. First, (α, β, γ) is minimal, and so neither $G_\gamma(\alpha)$ nor $G_\gamma(\beta)$ is a Type III graph. Second, the weight $w_\gamma(b) = 0$ in $G_\alpha(\gamma)$, which is a Type III graph. This implies that α and γ can be oriented so that all points of intersection occur are coherently oriented; in other words, (α, γ) is a *positive Heegaard diagram*. In this case, we consider $G_\gamma(\alpha)$.

Lemma 4.1 *The graph $G_\gamma(\alpha)$ is a square graph.*

Proof Since $G_\gamma(\alpha)$ does not contain a wave by assumption, we know that it falls under Type I or Type II. Refer to Figure 6. In addition, (α, γ) is a positive Heegaard diagram, so all cross-seams in $\Sigma_\gamma(\alpha)$ originating from γ_1^+ and terminating in $\gamma_2^\pm$ must terminate in the same vertex, say γ_2^+. Thus, $G_\gamma(\alpha)$ contains an edge connecting γ_1^+ to γ_2^+, its symmetric image connecting γ_1^- to γ_2^-, and no edges connecting $\gamma_1^\pm$ to $\gamma_2^\mp$. It follows that if $G_\gamma(\alpha)$ is a graph of Type II, then $w_\alpha(b) = 0$ and $G_\gamma(\alpha)$ contains a wave, a contradiction. Thus, $G_\gamma(\alpha)$ is a graph of Type I and $w_\alpha(b) = 0$, as desired. □

Recall that we assume that the edges in a square graph or a c-graph have nonzero weights. By Lemma 2.7, a graph for S^3 with only two edges is trivial, and such a graph must be Type I with $w(a) = w(b) = 0$ and $w(c) = w(d) = 1$. In keeping with Figure 6, we will refer to the edges $a_\pm$ of a square graph as *horizontal edges* and the edges c and d as *vertical edges*. Note that if (δ, ϵ) is a positive Heegaard diagram for S^3 such that $G_\delta(\epsilon)$ is a square graph, then $G_\delta(\epsilon)$ is not a Type III graph and as such, there is no wave based in δ for (δ, ϵ). By Theorem 2.6, it follows that there must be a wave based in ϵ. In this setting, we may keep track of the structure of $G_\delta(\epsilon)$ after doing wave surgery.

Lemma 4.2 *Suppose that (δ, ϵ) is a positive Heegaard diagram for S^3 such that $G_\delta(\epsilon)$ is a square graph, and let (δ, ϵ') be the result of doing wave surgery on (δ, ϵ). Then $G_\delta(\epsilon')$ is either a square graph or a c-graph.*

Proof Let $\delta = \{\delta_1, \delta_2\}$, $\epsilon = \{\epsilon_1, \epsilon_2\}$ and $\epsilon' = \{\epsilon_2, \epsilon_3\}$. Observe that ϵ_3 can be embedded disjointly from ϵ and consider $G_\delta(\epsilon_1 \cup \epsilon_2 \cup \epsilon_3)$. The curve ϵ_3 may introduce at most one additional edge; however, doing wave surgery introduces no new points of intersection and does not change the sign of existing intersection points. Therefore, (δ, ϵ') remains a positive diagram, and $G_\delta(\epsilon')$ cannot have such an edge. Thus $G_\delta(\epsilon')$ is a subgraph of $G_\delta(\epsilon)$.

We must show that $G_\delta(\epsilon')$ has more than two edges. Suppose by way of contradiction that $G_\delta(\epsilon')$ has two edges, so that it is the trivial diagram. It follows that $\iota(\delta, \epsilon_2) = 1$, so we may assume that $\iota(\delta_1, \epsilon_2) = 0$. By Lemma 3.5, (δ, ϵ) is not a Heegaard diagram for S^3, a contradiction. □

As an immediate consequence, we have the next lemma.

Lemma 4.3 *Suppose that (δ, ϵ) is a positive Heegaard diagram for S^3 and $G_\delta(\epsilon)$ is a square graph. Then there is a sequence of diagrams $(\delta, \epsilon) = (\delta, \epsilon^l), (\delta, \epsilon^{l-1}), \dots, (\delta, \epsilon^0)$ such that (δ, ϵ^i) is the result of wave surgery on (δ, ϵ^{i+1}), the graph $G_\delta(\epsilon^i)$ is a square graph for $i > 0$, and $G_\delta(\epsilon^0)$ is a c-graph.*

Proof By Lemma 4.2, $G_\delta(\epsilon^{l-1})$ is either a square graph or a c-graph. By induction on $(l-i)$, suppose that $G_\delta(\epsilon^i)$ is square graph. Observe that (δ, ϵ^i) always contains a wave, that the wave must be based in ϵ^i, and that δ is preserved by the wave surgery. Since $\iota(\delta, \epsilon^{i-1}) < \iota(\delta, \epsilon^i)$, finitely many applications of Lemma 4.2 yield the desired result. □

Observe that $\Sigma \setminus \nu(\delta \cup \epsilon^i)$, for $i > 0$, is a collection of rectangles and two octagons whose sides alternate between arcs in δ and arcs in ϵ^i. It follows that $G_{\epsilon^i}(\delta)$ is a Type III graph with $w_\delta(b) = 0$ and all other edge weights nonzero. As such, there is a unique way to do wave surgery on ϵ^i, and so we call the sequence $\{\epsilon = \epsilon^l, \epsilon^{l-1}, \dots, \epsilon^0\}$ guaranteed by Lemma 4.3 the *reducing sequence* for ϵ (without ambiguity). See Figure 14.

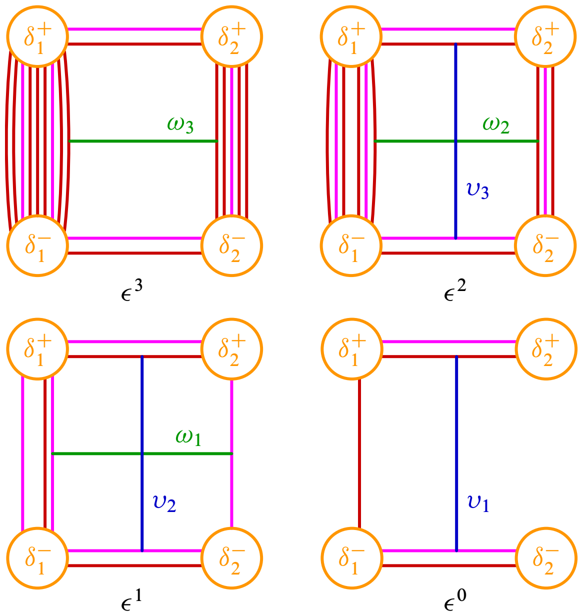

Figure 14: The reducing sequence for a diagram (δ, ϵ). All four graphs shown have alternating adjacency. Note that we have suppressed the gluing information needed to pass from $\Sigma_\delta(\epsilon^i)$ to a Heegaard diagram for S^3, although determining it is straightforward.

The involution of Σ maps one of the octagons cut out by δ and ϵ^i to the other octagon, and thus it suffices to examine one of the octagons, which we will call O_i. Since the diagram is positive, any wave or band for (δ, ϵ^i) must connect opposite arcs of ϵ^i contained in the boundary of O_i. We will call a wave ω parallel to a horizontal edge of $G_\delta(\epsilon^i)$ a *horizontal wave*; otherwise ω is parallel to a vertical edge and we call ω a *vertical wave*. The notions of a *horizontal band* and *vertical band* are defined similarly.

Let ω_i denote the wave for (δ, ϵ^i) such that surgery on ω_i yields (δ, ϵ^{i-1}), and let υ_i denote the band dual to each ω_i (so that banding (δ, ϵ^{i-1}) along υ_i yields (δ, ϵ^i)). Let $\epsilon^i = \{\epsilon_1^i, \epsilon_2^i\}$. Note that ω_i necessarily has endpoints in the same component of ϵ^i, whereas υ_{i+1} always connects ϵ_1^i to ϵ_2^i. In addition, υ_i is vertical if and only if ω_i is horizontal.

Lemma 4.4 *Every wave ω_i is a horizontal wave.*

Proof The band υ_1 is contained in the c-graph $G_\delta(\epsilon^0)$, and thus it must be a vertical band. It follows that ω_1 is a horizontal wave. Suppose by way of induction that ω_i is a horizontal wave in $G_\delta(\epsilon^i)$. Since the band υ_i is also contained in $G_\delta(\epsilon^i)$, we have that υ_{i+1} cannot be parallel to ω_i since ω_i connects the same component of ϵ^i while υ_{i+1} connects distinct components of ϵ^i. Thus, υ_{i+1} is vertical, which implies that its dual wave ω_{i+1} is horizontal as well. □

Here we introduce some terminology to better understand what happens to arcs of $\Sigma_\delta(\epsilon^i)$ when two curves are banded together. We say that two arcs in $\Sigma_\delta(\epsilon^i)$ are *adjacent* if they cobound a rectangle in $\Sigma \setminus \nu(\delta \cup \epsilon^i)$ with arcs in δ.

Lemma 4.5 *If ϵ_3^i is the result of banding ϵ_1^i to ϵ_2^i along υ_{i+1}, then $\iota(\delta, \epsilon_3^i) = \iota(\delta, \epsilon^i)$. In addition, in $\Sigma_\delta(\epsilon^i \cup \epsilon_3^i)$, every arc of ϵ^i is adjacent to an arc of ϵ_3^i.*

Proof In order to construct ϵ_3^i from ϵ_1^i and ϵ_2^i in Σ_δ, we begin with two parallel copies of ϵ_1^i and ϵ_2^i which have been pushed off of ϵ^i so that two of their horizontal arcs lie in the octagon O_i. Attaching the band υ_{i+1} to these two arcs transforms them into two vertical arcs and leaves all other arcs of ϵ^i unchanged, proving the first statement. Since the vertical band υ_{i+1} connects ϵ_1^i to ϵ_2^i, we have $w_{\epsilon_j^i}(a) \geq 1$ for $j = 1, 2$. Thus $w_{\epsilon^i}(a) \geq 2$, and every arc of ϵ^i is adjacent to some arc of its push-off that has not been changed by attaching the band υ_{i+1}. See Figure 15. □

Now, we return to the trisection diagram (α, β, γ). By Lemma 4.1, $G_\gamma(\alpha)$ is a square graph, and, as such, α has a reducing sequence $\{\alpha = \alpha^l, \alpha^{l-1}, \dots, \alpha^0\}$ by Lemma 4.3. The next lemma examines this reducing sequence more closely.

Lemma 4.6 *Suppose that α^0 and α^l have a curve α_*^0 in common. Then (α, β, γ) is not minimal.*

Proof Let $\alpha^l = \{\alpha_1^l, \alpha_*^0\}$. If (α, β, γ) is minimal, then $\iota(\alpha_1, \gamma) \geq \iota(\alpha_2, \gamma)$, and by Lemma 4.5, $\iota(\alpha_1^l, \gamma) \geq \iota(\alpha^0, \gamma) > \iota(\alpha_*^0, \gamma)$. It follows that $\alpha_1 = \alpha_1^l$ and $\alpha_2 = \alpha_*^0$. In addition, we have that $G_\gamma(\alpha_2)$ is a subgraph of the c-graph $G_\gamma(\alpha^0)$. If $G_\gamma(\alpha_2)$ contains a cross-seam, them (α, β, γ) is not minimal by Lemma 3.3. Otherwise, $\iota(\alpha_2, \gamma_j) = 0$ for either $j = 1$ or $j = 2$; therefore, (α, γ) is not a Heegaard diagram for S^3 by Lemma 3.5. □

In order to deal with the intricacy of the remaining case in Section 5, we will need to keep track of the cyclic ordering of arcs at the vertices of $\Sigma_\alpha(\gamma)$. One way to obtain this information is to understand adjacency in $\Sigma_\gamma(\alpha)$, since the regions of $\Sigma \setminus \nu(\alpha \cup \gamma)$ appear in both graphs. We say that (α^i, γ) has *upper alternating adjacency* (resp. *lower alternating adjacency*) if the horizontal boundary arc of O_i in the edge a^+ (resp. a^-) is adjacent to an arc in the opposite curve of α^i. See Figure 14. If (α^i, γ) has both upper and lower alternating adjacencies, we say that it has *alternating adjacency*. The remainder of the section includes several lemmas which establish that $G_\gamma(\alpha)$ has alternating adjacency.

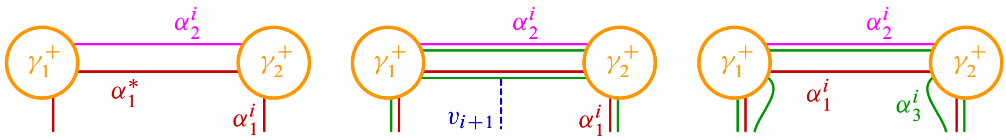

Figure 15: The graph $G_\gamma(\alpha^i)$: (left) exhibits upper alternating adjacency, (middle) shows the parallel push offs of α_1^i and α_2^i, and (right) shows α_3^i, the result of the banding.

Lemma 4.7 *If (α^i,γ) has upper (resp. lower) alternating adjacency, then (α^{i+1},γ) has upper (resp. lower) alternating adjacency.*

Proof By Lemma 4.4, the band υ_{i+1} dual to the horizontal wave ω_{i+1} is vertical. Let α_3^i correspond to the curve created by banding α_1^i to α_2^i along υ_{i+1}, so that α^{i+1} is either $\{\alpha_1^i,\alpha_3^i\}$ or $\{\alpha_2^i,\alpha_3^i\}$.

Let $\alpha^*=\{\alpha_1^i,\alpha_2^i,\alpha_3^i\}$ and consider $\Sigma_\gamma(\alpha^*)$, letting O_* denote the octagon in $\Sigma\setminus\nu(\alpha^*\cup\gamma)$ such that $O_*\subset O_{i+1}$. By the proof of Lemma 4.5, both vertical arcs of ∂O_* in α^* are in α_3^i, and ∂O_* has one horizontal arc in each of α_1^i and α_2^i; call these arcs α_1^* and α_2^*, respectively. Suppose that α_1^* is in the edge a_+ (resp. a_-). Then by Lemma 4.5, α_1^* is adjacent to an arc in α_3^i, which in turn is adjacent to an arc in α_2^i by the upper (resp. lower) alternating adjacency of $\{\alpha^i,\gamma\}$. Thus, independently of whether $\alpha_1^i\in\alpha^{i+1}$ or $\alpha_2^i\in\alpha^{i+1}$, the diagram (α^{i+1},γ) must also have upper (resp. lower) alternating adjacency. See Figure 15. □

Note that, in the proof of Lemma 4.7, if $\alpha_1^i\in\alpha^{i+1}$, then we need not require that (α^i,γ) have upper (resp. lower) alternating adjacency. The next lemma follows immediately.

Lemma 4.8 *If an arc in α_j^i is in the upper (resp. lower) horizontal boundary of O_i and $\alpha_j^i\in\alpha^{i+1}$, then (α^{i+1},γ) has upper (resp. lower) alternating adjacency.*

Now, we combine Lemmas 4.7 and 4.8 to determine adjacency for (α,γ) when (α,β,γ) is minimal.

Lemma 4.9 *If α^0 and α^l do not have a curve in common, then (α^l,γ) has alternating adjacency.*

Proof Note that the band υ_1 is vertical, so it connects α_1^0 to α_2^0. Suppose without loss of generality that υ_1 connects an upper horizontal arc of α_1^0 to a lower horizontal arc of α_2^0, and suppose further that $\alpha_1^0\in\alpha^1$. By Lemma 4.8, (α^1,γ) has upper alternating adjacency, and by Lemma 4.7, (α^i,γ) has upper alternating adjacency for all $i > 0$. Let k be the largest index such that $\alpha_1^0 \in \alpha^k$, and suppose $\alpha^k = \{\alpha_1^0, \alpha_2^k\}$. By assumption, $k < l$, and following the proof of Lemma 4.7, we have that the upper horizontal boundary of O_k is an arc in α_1^0 while the lower boundary is an arc in α_2^k. Since $\alpha_1^0 \notin \alpha^{k+1}$, we have $\alpha_2^k \in \alpha^{k+1}$ and thus (α^{k+1}, γ) has lower alternating adjacency by Lemma 4.8. Finally, Lemma 4.7 implies that (α^i, γ) has lower alternating adjacency for all $i > k$, and thus (α_l, γ) has alternating adjacency as desired. $\square$

In summary, we have the following.

Lemma 4.10 *If $G_\gamma(\alpha)$ is a square graph, then (α, γ) has alternating adjacency or (α, β, γ) is not minimal.*

Proof By Lemma 4.3, there is a reducing sequence $\alpha = \alpha^l, \alpha^{l-1}, \dots, \alpha^0$ for (α, γ). If α^l and α^0 have no curve in common, then (α, γ) has alternating adjacency by Lemma 4.9. Otherwise, (α, β, γ) is not minimal by Lemma 4.6. $\square$

5 The case in which $m = -n \pm 1$

Up to this point in the paper, our arguments have not required information about the cyclic ordering of the intersections of arcs with the fat vertices in $\Sigma_\alpha(\beta, \gamma)$. However, this case is significantly more delicate than the preceding ones, and thus we must be more precise about how we recover (α, β, γ) from $\Sigma_\alpha(\beta, \gamma)$. We will assume that β has been isotoped so that all inessential intersections between β and γ correspond to triangles based in α_1^+ or α_2^+.

Let us recall our current assumptions. First, neither diagram (α, γ) or (β, γ) is trivial, there is no wave based in γ, and the trisection diagram (α, β, γ) is minimal. Second, $G_\alpha(\gamma)$ is a Type III graph such that $w_b(\gamma) = 0$ and all other weights are nonzero. This implies that $\Sigma \setminus \nu(\alpha \cup \gamma)$ contains some number of rectangular regions and two octagonal regions, which we call O' and O''. Third, $G_\gamma(\alpha)$ is a square graph with alternating adjacency and (α, γ) contains a wave based in α_1. Fourth, by Lemma 3.6 we may assume that (β, γ) does not contain a wave based in γ, so that it must contain a wave based in β. Finally, the slopes of the β–arcs in $\Sigma_\alpha(\beta, \gamma)$ are both $-(n \pm 1)/n$. Since the slopes of the β–arcs differ from the slopes of the γ–arcs, and since there are no inessential intersection points based at α_i^-, we have that β_i intersects α_i^- in an edge of either O' or O'' for $i = 1, 2$.

Lemma 5.1 *If $G_\gamma(\alpha)$ has an edge of weight one, then (α, β, γ) is not minimal.*

Proof Suppose that in the square graph $G_\gamma(\alpha)$, one of the edges $a_\pm$, c or d has weight one (see Figure 6). Note that each octagon O' and O'' contains three boundary components in α_1 and one boundary component in α_2. By Lemma 4.3, there is a horizontal wave contained in $\Sigma_\gamma(\alpha)$, and so both vertical boundary arcs of O' in $\Sigma_\gamma(\alpha)$ are in α_1. This means that O' has horizontal boundary arcs in both α_1 and α_2; thus $w_\alpha(a) \geq 2$.

Consequently, we suppose without loss of generality that $w_\alpha(d) = 1$. If $w_{\alpha_2}(d) = 0$, then $G_\gamma(\alpha_2)$ is a Type III graph containing a cross-seam and (α, β, γ) is not minimal by Lemma 3.3. If $w_{\alpha_2}(d) = 1$, let $\gamma_3 = \Delta_a$ and let $\gamma' = \{\gamma_2, \gamma_3\}$. Then γ' is related to γ by a handle slide, so by the minimality of (α, β, γ), we have that $\iota(\alpha_2, \gamma) \leq \iota(\alpha_2, \gamma')$ and thus $\iota(\alpha_2, \gamma_1) \leq \iota(\alpha_2, \gamma_3)$. But this implies that $w_{\alpha_2}(a) + w_{\alpha_2}(c) \leq w_{\alpha_2}(c) + 1$ and so $w_{\alpha_2}(a) = 1$ and $\iota(\alpha_2, \gamma) = \iota(\alpha_2, \gamma')$. By the above arguments, $w_{\alpha_1}(a) > 0$, and so

$$\iota(\alpha_1, \gamma') = w_{\alpha_1}(a) + w_{\alpha_1}(c) < 2w_{\alpha_1}(a) + w_{\alpha_1}(c) = \iota(\alpha_1, \gamma).$$

This implies $\mathcal{C}(\alpha, \beta, \gamma') < \mathcal{C}(\alpha, \beta, \gamma)$, as desired. □

The next lemma uses our work from Section 4 to determine how β_1 meets α_1^+ relative to γ in $\Sigma_\alpha(\beta, \gamma)$.

Lemma 5.2 *For $i = 1, 2$, the arc β_i intersects α_i^+ between parallel arcs of γ in the edge a_+.*

Proof Let $\alpha_* \subset \alpha_1$ denote the boundary arc of the octagon, say O', which contains the intersection of β_1 with α_1^-, and observe by examination of $\Sigma_\alpha(\gamma)$ that in O' the arc α_* is opposite an arc in α_2. See Figure 16. By Lemma 4.4, the wave ω for (α, γ) is horizontal, and thus opposite pairs of vertical arcs in $\Sigma_\gamma(\alpha)$ contained in $\partial O'$ meet ω and are in the same curve α_1. This implies that α_* is a horizontal arc in one of $a_\pm$ in $G_\gamma(\alpha)$. Now, we invoke Lemma 4.10, which asserts that (α, γ) has alternating adjacency. It follows that α_* also cobounds a rectangular component R of $\Sigma \setminus \nu(\alpha \cup \gamma)$, and is opposite an arc of α_2 in R. Since the only such rectangles in $\Sigma_\alpha(\gamma)$ lie between parallel arcs in the edge a_+, the desired statement follows. A parallel argument shows that the same is true for β_2 and α_2^+. □

We note that if the weight of the edge a_+ in $\Sigma_\gamma(\alpha)$ is large, then β_1 and β_2 will not necessarily intersect α_1^+ and α_2^+ between the *same* set of parallel arcs in the edge a_+ in $\Sigma_\alpha(\gamma)$ as in Figure 16. We will also keep track of how β and γ intersect near the vertices α_i^+, but we require the following lemma first.

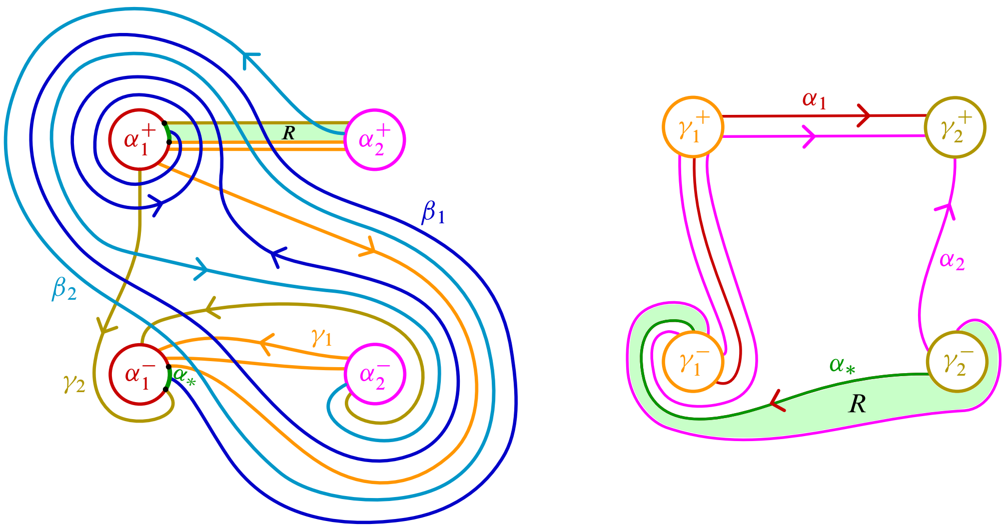

Figure 16: An example of $\Sigma_\alpha(\beta,\gamma)$ (left), along with the corresponding $\Sigma_\gamma(\alpha)$ (right). At left, $W_1 = \frac{13}{5}$, $W_2 = \frac{1}{3}$, and the slope of the β–curves is $-\frac{3}{4}$. Notice that (β,γ) is a positive Heegaard diagram according with Lemma 5.5. We have assumed that there are no inessential intersections at α_1^-, which prescribes that β_1 leaves this vertex as shown. The existence of the rectangle R on the right dictates that β_1 intersects α_1^+ as shown, as in Lemma 5.2.

Lemma 5.3 *All essential points of $\beta \cap \gamma$ are coherently oriented.*

Proof Consider $\Sigma_\alpha(\beta,\gamma)$. Choose an orientation for γ and note that since (α,γ) is a positive diagram, each arc corresponding to a single edge in $G_\alpha(\gamma)$ has the same orientation. Thus, we may orient the edges of $G_\alpha(\gamma)$ so that oriented arcs originate in α_1^+ and α_2^- and terminate in α_1^- and α_2^+. Now, we may orient β_1 and β_2 so that all essential points of intersection of β and the edge a_+ are coherently oriented. The involution J preserves β_1 and β_2 setwise while reversing their orientations; hence essential points of $\beta \cap a_+$ and $\beta \cap a_-$ are coherently oriented. Since the slope $\frac{m}{n}$ of β is negative, it follows that essential points of $\beta \cap c$ and $\beta \cap a_-$ are coherently oriented as well. Finally, the homeomorphism ψ of $G_\alpha(\beta,\gamma)$ swaps γ–arcs in the edge c with those in the edge d while preserving their orientations and maps β–arcs of slope $-(n+1)/n$ to arcs of slope $-(n+1)/(n+2)$. Thus, by applying ψ to $G_\alpha(\beta,\gamma)$, we may make a parallel argument that all essential points of $\beta \cap d$ and $\beta \cap a_-$ are coherently oriented. We conclude that all essential intersections of β and γ in $G_\alpha(\beta,\gamma)$ are coherently oriented. See Figure 16 and Figure 18 (left). □

Now, we define the *winding of β_j relative to γ at α_i*, denoted $W_i(\beta_j)$ (or simply W_i when the relevant β–curve is clear), to be the number of inessential intersections (counted with sign) of β and γ that are vertices of triangles based in α_i^+ divided by $\iota(\alpha_i,\gamma)$ for $i=1,2$. An inessential intersection is positive if it is coherently oriented with the essential points of intersection and negative otherwise. In the next four lemmas, we place restrictions on the values of W_1 and W_2 by showing that certain combinations of windings imply that (α,β,γ) is not minimal or (β,γ) does not contain a wave based in β, which contradicts our prior assumptions.

Lemma 5.4 *If W_1 or W_2 is an integer, then (α,β,γ) is not minimal.*

Proof If the winding of β relative to γ at α_1 is an integer, then the two octagonal regions of $\Sigma \setminus \nu(\alpha\cup\gamma)$ are separated by a single arc contained in α_1. It follows that in the square graph $G_\gamma(\alpha)$, one of the edges has weight one and thus (α,β,γ) is not minimal by Lemma 5.1. □

Lemma 5.5 *If $W_i<0$ for either $i=1$ or $i=2$, then $G_\beta(\gamma)$ does not contain a wave.*

Proof Consider $\Sigma_\alpha(\beta,\gamma)$ and note that there is a subarc of a γ–arc contained in the edge a_+ with endpoints in essential intersections of β_1 and β_2; thus $G_\beta(\gamma)$ contains an edge connecting β_1^+ to β_2^+. For any arc of γ originating from α_1^+, its first essential intersection is with β_2. Suppose that $\beta_1\cap\gamma$ contains an inessential point of intersection based in α_1^+ which is oriented opposite the essential points of intersection. Then there is an arc of γ connecting an inessential point of intersection with β_1 with an oppositely oriented essential point of intersection with β_2. It follows that $G_\beta(\gamma)$ contains an edge connecting β_1^+ to β_2^-, and by the argument of Lemma 2.9, $G_\beta(\gamma)$ does not contain a wave.

A similar argument shows that if $\beta_2\cap\gamma$ has an oppositely oriented inessential intersection point based in α_2^+, then $G_\beta(\gamma)$ again does not contain a wave. □

It follows that $W_i>0$ for $i=1,2$; hence, (β,γ) is a positive diagram, and $G_\gamma(\beta)$ is a square graph.

Lemma 5.6 *If $0<W_i<1$ for both $i=1,2$, then $G_\beta(\gamma)$ does not contain a wave.*

Proof If $0<W_1<1$, then there is an arc γ' in γ whose endpoints lie in α_1^+ and β_2. See Figure 17 (left). By assumption, there are no inessential intersections of β and γ based at α_1^-, and so every arc of γ with an endpoint on α_1^- meets β_2. Thus, the

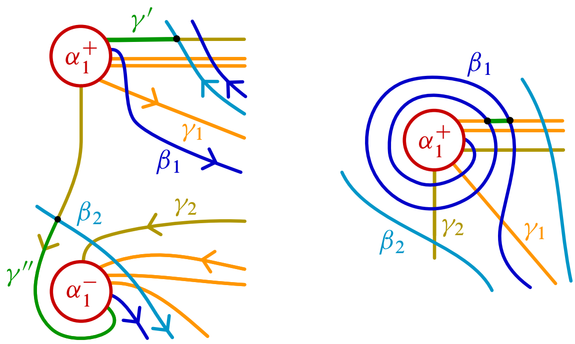

Figure 17: The cases in which $W_1 < 1$ (left) and $W_1 > 1$ (right)

endpoint of γ' in α_1^+ is the endpoint of another arc γ'' connecting α_1^- to β_2, and in the graph $G_\beta(\gamma)$, the arc $\gamma' \cup \gamma''$ is in an edge connecting β_2^+ to β_2^-.

A parallel argument shows that if $0 < W_2 < 1$, then $G_\beta(\gamma)$ contains an arc in an edge connecting β_1^+ to β_1^-. Together, these two arcs constitute a wave-busting pair, and by Lemma 2.9, we have that $G_\beta(\gamma)$ does not contain a wave. $\square$

Lemma 5.7 *If $W_i > 1$ for both $i = 1, 2$, then $G_\beta(\gamma)$ does not contain a wave.*

Proof If $W_1 > 1$, then there is an arc of γ emanating from α_1^+ that meets β in two consecutive inessential points of $\beta_1 \cap \gamma$ which are coherently oriented. See Figure 17 (right). Such an arc contributes an edge connecting β_1^+ to β_1^- in $G_\beta(\gamma)$. Similarly, if $W_2 > 1$, then $G_\beta(\gamma)$ contains an edge connecting β_2^+ and β_2^-, and thus $G_\beta(\gamma)$ contains a wave-busting pair. We conclude that $G_\beta(\gamma)$ does not contain a wave by Lemma 2.9. $\square$

Thus, we are left with two cases: $W_i > 1$ and $0 < W_j < 1$ for $\{i, j\} = \{1, 2\}$. Unfortunately, in both of these cases it is possible for (β, γ) to admit a wave based in β, and so a more sophisticated argument is required—especially since the result β' of wave surgery on β no longer satisfies $\iota(\alpha, \beta') = 2$. Hence, we must appeal to a line of reasoning other than that involving minimal complexity $\mathcal{C}(\alpha, \beta, \gamma)$. In the remaining two lemmas, we complete the proof of Proposition 3.1 by using Lemma 4.3 to determine that (β, γ) is not a Heegaard diagram for S^3.

Lemma 5.8 *If $W_1 > 1$ and $0 < W_2 < 1$, then (β, γ) is not a Heegaard diagram for S^3.*

Proof First, we note that (β, γ) contains a wave ω: Since $W_1 > 1$, an octagonal component O' of $\Sigma \setminus \nu(\alpha \cup \gamma)$ contains an arc β_* of β_1, both of whose endpoints

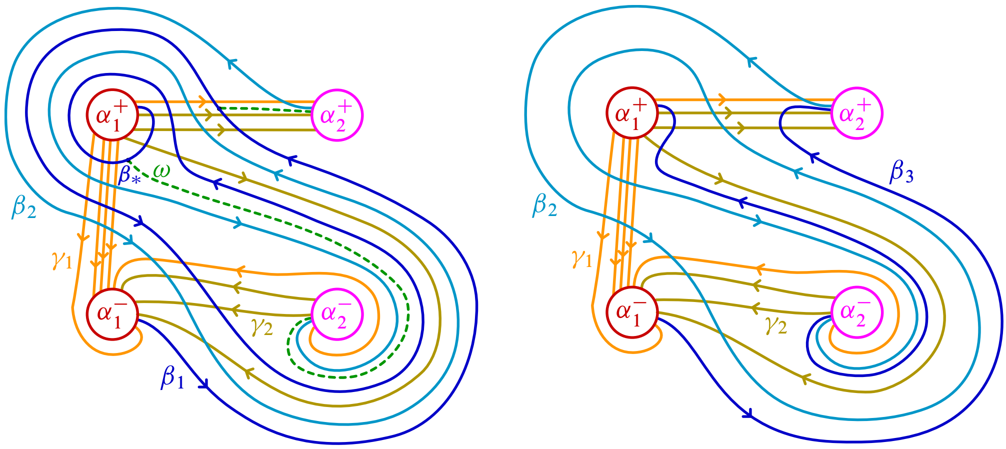

Figure 18: Before (left) and after (right) the surgery along the β_1–wave ω in $\Sigma \setminus \nu(\alpha \cup \gamma)$

are inessential points of $\beta \cap \gamma$, such that there is an arc ω which avoids $\beta \cup \gamma$ in its interior, has one endpoint on β_*, meets α_2^- in a single point contained in O', and continues through a rectangular component of $\Sigma \setminus \nu(\alpha \cup \gamma)$ to meet β_1 in an oppositely oriented point. See Figure 18 (left).

As depicted in Figure 18 (right), we let β_3 denote the new curve created by doing surgery on β_1 along ω. The curve β_3 intersects $G_\alpha(\beta, \gamma)$ in two arcs, each having slope $-\frac{1}{1}$. Further, $W_1(\beta_3) = W_1(\beta_1) - 1$ while $W_2(\beta_3) = 1 - W_2(\beta_2)$. Let $\beta' = \{\beta_2, \beta_3\}$. We claim that (β', γ) does not contain a wave based in γ. By Lemma 5.1, each edge c and d of $G_\gamma(\alpha)$ has weight at least two, which implies that $\Sigma \setminus \nu(\alpha \cup \gamma)$ contains two rectangles R_1 and R_2 such that R_i has two opposite boundary components contained in γ_i. See Figure 19. Since the slope of β_2 is $-(n \pm 1)/n$, we have that β_2 intersects each edge of $G_\alpha(\gamma)$ essentially, and thus β_2 meets R_1 and R_2 in arcs β_*^1 and β_*^2. In $G_\gamma(\beta')$, the pair $\{\beta_*^1, \beta_*^2\}$ is a wave-busting pair, and thus $G_\gamma(\beta')$ does not contain a wave by Lemma 2.9.

Now, suppose that $W_1(\beta_1) > 2$, so that $W_1(\beta_3) > 1$. In this case, as in the proof of Lemma 5.7, there is an arc γ_3^* of γ which meets β_3 in its coherently oriented endpoints and avoids β' in its interior. In addition, $0 < W_2(\beta_2) < 1$; hence there is an arc γ_2^* contained in the edge a_+ which meets β_2 in its coherently oriented endpoints and avoids β' in its interior. See Figure 20. It follows that $\{\gamma_2^*, \gamma_3^*\}$ is a wave-busting pair and $G_{\beta'}(\gamma)$ does not contain a wave by Lemma 2.9. In this case, we conclude that (β', γ), and therefore (β, γ), is not a Heegaard diagram for S^3 by Theorem 2.6.

On the other hand, suppose that $1 < W_1(\beta_1) < 2$, so that $0 < W_1(\beta_3) < 1$. In this case the situation is considerably more complicated. Suppose (by applying ψ if necessary

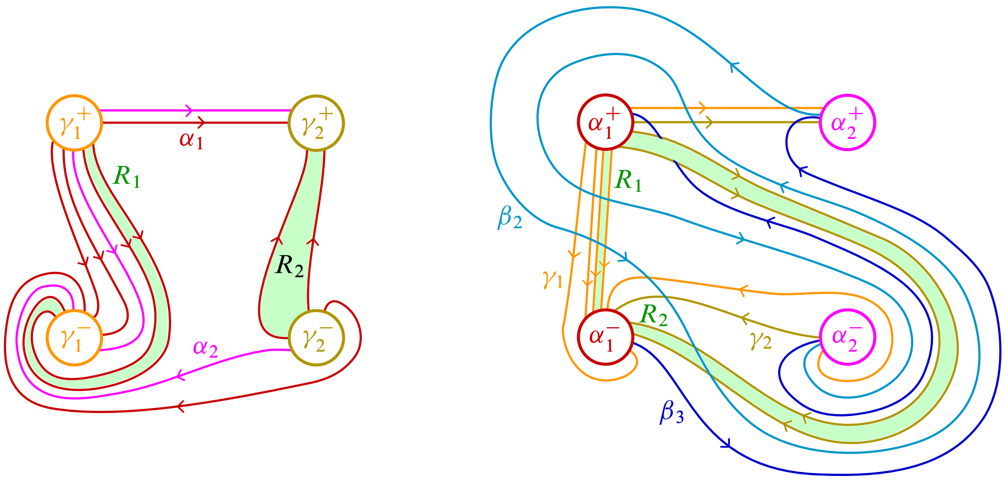

Figure 19: The fact that no edge in $G_\gamma(\alpha)$ has weight one gives rise to a pair of rectangles in $\Sigma_\gamma(\alpha)$ (left) that in turn produce a wave-busting pair (β_*^1, β_*^2) in $\Sigma_\alpha(\beta', \gamma)$ (right). Each arc β_*^i comes from a component of $\beta_2 \cap R_i$.

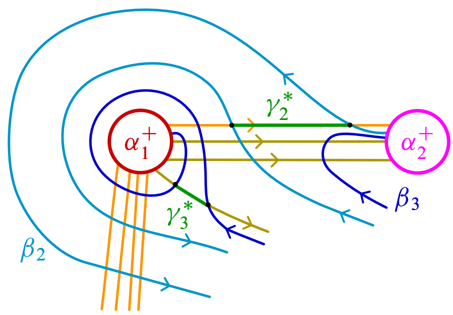

Figure 20: A local view of the α_i^+ in the case that $W_1(\beta_3) > 1$. In this case, $\{\gamma_2^*, \gamma_3^*\}$ is a wave-busting pair of arcs.

and noting that ψ does not alter the values of W_1 and W_2) that the slope of β_2 is $\frac{m}{n} = -(|n|-1)/|n|$. By Lemma 2.4, the curve β_2 has exactly $k = \frac{|n|}{2} - 1$ essential intersections with each arc in the edges $a_\pm$. Consider the arc ω' connecting α_1^+ to α_2^- in the octagon O'. By the symmetry of $G_\alpha(\beta, \gamma)$ under the involution J and Lemma 5.2, if we extend ω' through α_1, it intersects α_1^- between parallel arcs corresponding to the edge a_-. See Figure 21. Thus, ω' may be extended from α_1^- and from α_2^+ to an arc ω^* which avoids β_3 and meets β_2 in $2k$ points, where k of these points arise from the k essential intersections of β_2 with a_+ and the other k points arise from the intersections of β_2 with a_-. Moreover, each intersection point of the first

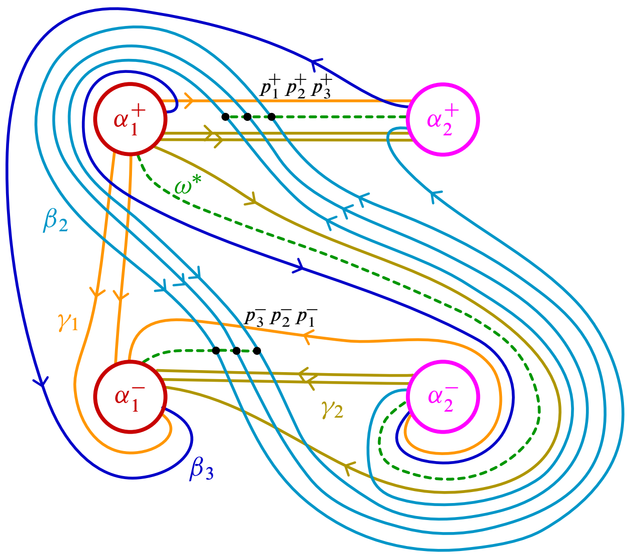

Figure 21: An example of $\Sigma_\alpha(\beta', \gamma)$ in the case that $W_1(\beta_3) < 1$, along with the nested wave sequence ω^*. The slope of β_2 is $-\frac{7}{6}$.

type is naturally paired with an intersection point of the second type. For this reason, we will label the points of intersection of ω^* with β_2 as $p_1^+, \dots, p_k^+, p_k^-, \dots, p_1^-$.

We call the arc ω^* a *nested wave sequence*, because it gives rise to k successive waves. Let ω_i denote the subarc of ω^* connecting p_i^+ to p_i^-, so that $\omega^* = \omega_1$. Let $\beta_k^* = \beta_2$ and let $\beta_k' = \beta' = \{\beta_k^*, \beta_3\}$. Observe that ω_k is a wave for (β_k', γ) based in β_k^*, and let $\beta_{k-1}' = \{\beta_{k-1}^*, \beta_3\}$ be the result of doing surgery on β_k^* along ω_k. Inductively, ω^* gives rise a to sequence of wave surgeries: We will let $\beta_i' = \{\beta_i^*, \beta_3\}$, where β_i^* is the result of doing surgery on β_{i+1}^* along the wave ω_{i+1} for the diagram β_{i+1}'. See Figure 22.

Observe that by construction, the curve β_0^* does not intersect the edge a_+ essentially; however, for $i \geq 1$, we have that β_i^* intersects all edges of $G_\alpha(\gamma)$ essentially. As in the argument for the above case (shown in Figure 19), it follows that for $i \geq 1$, β_i^* intersects the rectangles R_1 and R_2 in a wave-busting pairs; thus $G_\gamma(\beta_i')$ does not contain a wave by Lemma 2.9. Lemmas 2.8 and 4.2 then imply that $G_\gamma(\beta_i')$ is a square graph, and so $\beta, \beta_k', \dots, \beta_1'$ are contained in the reducing sequence for β with respect to γ guaranteed by Lemma 4.3. Therefore, $G_\gamma(\beta_0')$ is either a square graph or a c-graph and as such (β_0', γ) contains a wave based in β_0'.

Now, each arc of β_0^* has nonzero winding relative to γ at both α_1^- and α_2^+, and so there is an arc γ_0^* that originates at the innermost inessential point of $\beta_0^* \cap \gamma$ based at α_2^+,

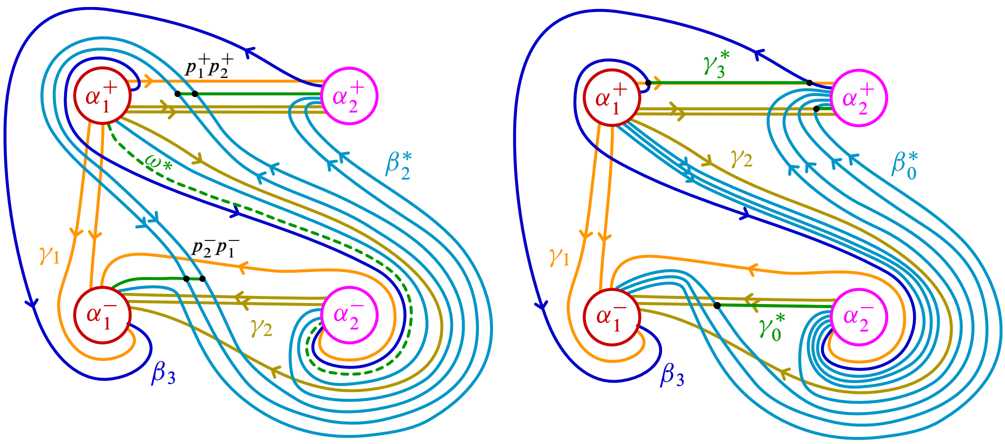

Figure 22: The fat graph $\Sigma_\alpha(\beta', \gamma)$ after the first wave move in ω^* (left) and after all of the wave moves in ω^* (right). On the right, we see a pair of wave-busting arcs $\{\gamma_0^*, \gamma_3^*\}$.

travels through α_2, and extends from α_2^- along an arc in the edge a_- to meet β_0^* and has coherently oriented endpoints. See Figure 22 (right). In addition, by Lemma 5.2, the curve β_3 meets α_1^+ in a point between parallel arcs in the edge a_+, and we have already established that $0 < W_i(\beta_3) < 1$ for $i = 1, 2$. Thus, there is an arc γ_3^* in the edge a_+ that meets β_3 in its coherently oriented endpoints. The pair (γ_0^*, γ_3^*) is a wave-busting pair for $G_{\beta_0'}(\gamma)$, implying that (β_0', γ) does not contain a wave based in β_0' by Lemma 2.9, contradicting our previous assumptions. This completes the proof of the lemma. $\square$

There is one final case to consider, which mirrors the proof of Lemma 5.8 almost exactly, and so we give an abbreviated proof.

Lemma 5.9 *If $0 < W_1 < 1$ and $W_2 > 1$, then (β, γ) is not a Heegaard diagram for S^3.*

Proof The proof in this case is virtually identical to the proof of Lemma 5.8. As above, (β, γ) contains a wave based in β: An arc connecting α_1^+ to α_2^- extends to a wave ω based in β_2. Let $\beta' = \{\beta_1, \beta_3\}$ be the result of surgery on β along ω. If $W_2 > 2$, then (β', γ) and thus (β, γ) is not a Heegaard diagram for S^3. Otherwise, we find a nested wave sequence that gives rise to a partial reducing sequence $\beta, \beta_k', \dots, \beta_1'$ for β with respect to γ, implying that (β_0', γ) contains a wave based in β_0', where $\beta_0' = \{\beta_0^*, \beta_3\}$. The only distinction in this proof is that the roles of γ_0^* and γ_3^* from the proof of Lemma 5.8 are reversed. Now, γ_0^* is an arc in a_+ connecting coherently oriented points of β_0^* whereas γ_3^* runs over α_2 to connect coherently points of β_3. See Figure 23. Of course, $\{\gamma_0^*, \gamma_3^*\}$ remains a wave-busting pair for $G_{\beta_0'}(\gamma)$, a contradiction. $\square$

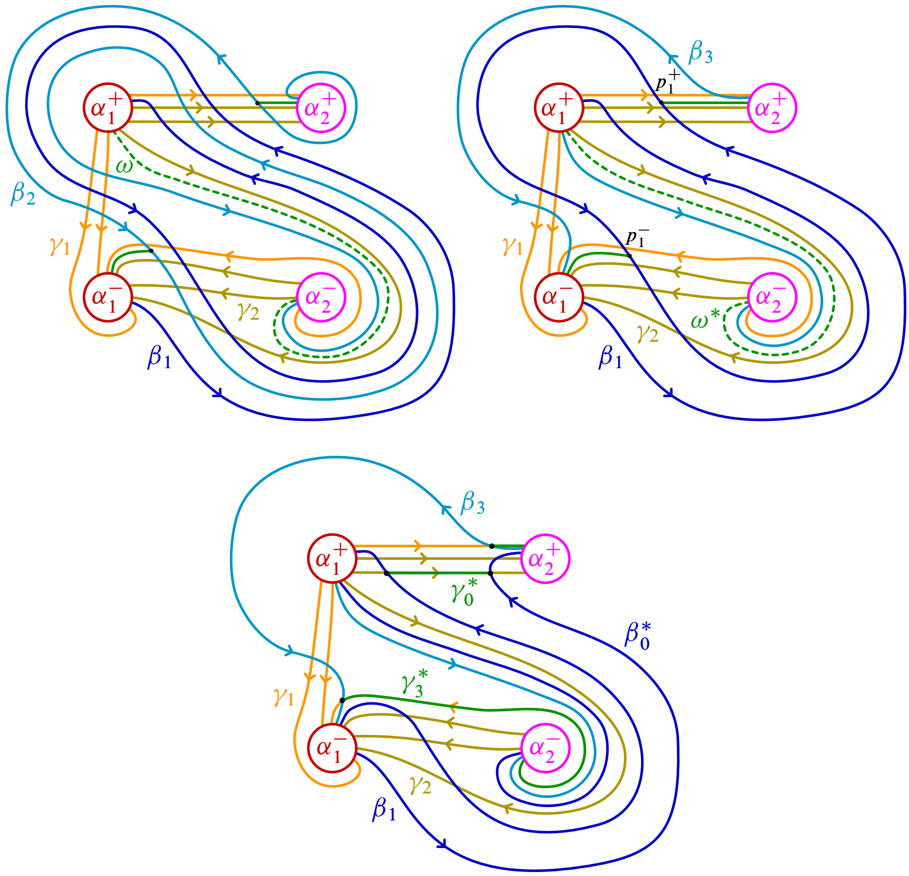

Figure 23: The case in which $0 < W_1 < 1$ and $W_2 > 1$. The wave ω (top left), the nested wave sequence ω^* (top right), and the wave-busting pair $\{\gamma_0^*, \gamma_3^*\}$ (bottom) appear in the proof of Lemma 5.9.

We summarize the previous three sections.

Proof of Proposition 3.1 Suppose that (α, β, γ) is a minimal diagram for the genus-two trisection $X = X_1 \cup X_2 \cup X_3$ which satisfies $\iota(\alpha, \beta) = 2$, and suppose by way of contradiction that (α, γ) is not the trivial diagram. By the minimality of (α, β, γ), the diagram (β, γ) is also nontrivial. By Theorem 2.6, there are waves for both (α, γ) and (β, γ). By Lemma 3.2, (α, γ) does not contain a wave based in γ, and so $G_\alpha(\gamma)$ is a Type III graph, and we parametrize the two arcs of β in $\Sigma_\alpha(\beta, \gamma)$ with slope $\frac{m}{n}$. By Lemma 3.4, we have $m \neq 1$. In addition, Lemma 3.6 asserts (β, γ) does not contain a wave based in γ; therefore, it must have a wave based in β, and so $w_\gamma(b) = 0$ in $G_\alpha(\gamma)$ by Lemma 3.7.

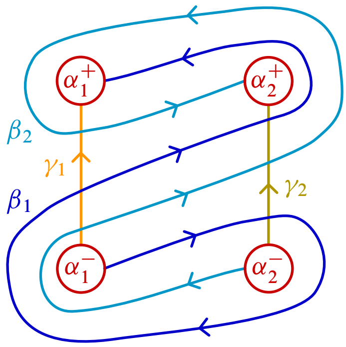

Figure 24: An example of $\Sigma_\alpha(\beta,\gamma)$ in which the β–arcs have slope $\frac{1}{4}$, and the winding numbers W_1 and W_2 are both zero

Lemma 3.8 then implies that $\frac{m}{n}$ must be $-(n\pm 1)/n$, the case dealt with in this section. Concerning the winding numbers W_1 and W_2 of β with respect to γ, Lemmas 5.4, 5.5, 5.6 and 5.7 yield that $W_i > 1$ and $0 < W_j < 1$ for $\{i,j\}=\{1,2\}$. However, Lemmas 5.8 and 5.9 show that, in these two cases, (β,γ) is not a Heegaard diagram for S^3, a contradiction. We conclude that (α,γ) must be the trivial diagram. □

6 The case in which $\iota(\alpha,\gamma)=2$

Proposition 3.1 reveals that a minimal trisection diagram (α,β,γ) for a $(2,0)$–trisection satisfies $\iota(\alpha,\beta)=\iota(\alpha,\gamma)=2$. In this section, we show that, in addition, $\iota(\beta,\gamma)$ must also equal 2, which will complete the proof of the main theorem. The strategy is to consider the graph $\Sigma_\alpha(\beta,\gamma)$, in which each of curve in $\beta\cup\gamma$ contributes a single arc, and the slopes of β–arcs agree and the slopes of the γ–arcs agree. We may suppose without loss of generality that the γ–arcs have slope $\frac{1}{0}$ and that the β–arcs have slope $\frac{m}{n}$, where m is odd and n is even. Although this choice is somewhat nonstandard (it seems more natural to let slopes of γ–arcs vary) this convention allows us to import techniques from Section 5 more consistently. We will orient β and γ so that oriented arcs in $\Sigma_\alpha(\beta,\gamma)$ originate in α_i^- and terminate in α_i^+, as shown in Figure 24.

Recall the homeomorphisms φ and φ^{-1} of Σ_α defined in Section 2 and given by

$$M_\varphi=\begin{bmatrix}1&1\\0&1\end{bmatrix}\quad\text{and}\quad M_{\varphi^{-1}}=\begin{bmatrix}1&-1\\0&1\end{bmatrix}.$$

Note that these homeomorphisms preserve γ setwise but exchange α_2^+ and α_2^-, which results in a reversal of orientations of β_2 and γ_2. When we apply such a homeomorphism, we will always relabel vertices and reorient the curves in β and γ to match the conventions given above. See Figure 25.

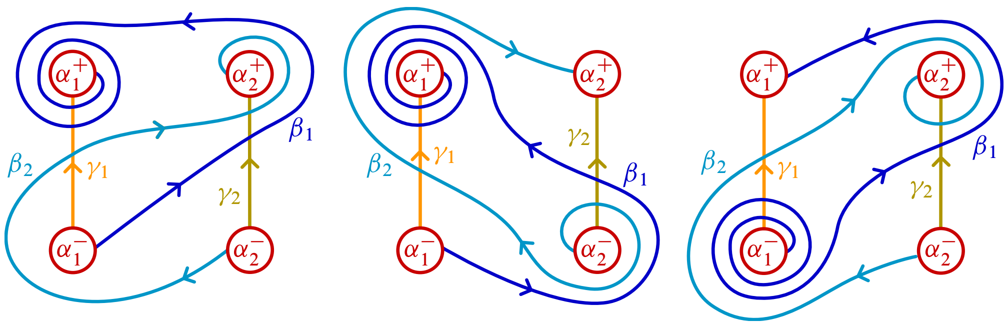

Figure 25: An example of $\Sigma_\alpha(\beta, \gamma)$ (left) and the effects of the homeomorphisms φ^{-1} (middle) and $H \circ \varphi^{-1}$ (right)

Lemma 6.1 *The slope $\frac{m}{n}$ for β may be chosen so that $-\frac{1}{2} < \frac{m}{n} \leq \frac{1}{2}$.*

Proof Observe that $\varphi^{\pm 1}$ preserves γ and takes the slope $\frac{m}{n}$ of β to $(m \pm n)/n$. Since there exists an integer m' such that $-\frac{n}{2} < m' \leq \frac{n}{2}$ and $m' \equiv m \bmod n$, we may map β–arcs to arcs of slope $\frac{m'}{n}$ by repeated applications of either φ or φ^{-1}, as desired. □

As in previous sections, we will systematically rule out various configurations of β and γ in Σ_α until the only remaining cases are the ones in which $\iota(\beta, \gamma) = 2$, shown in Figure 2.

Lemma 6.2 *If $m > 1$, then (β, γ) is not a Heegaard splitting of S^3.*

Proof We will show that there are wave-busting arcs contained in both β and γ, implying that (β, γ) does not contain a wave. By applying the horizontal reflection H if necessary, we may suppose without loss of generality that $-\frac{1}{2} < \frac{m}{n} < -\frac{1}{|n|}$. Then, $\frac{m}{n} \neq -(n \pm 1)/n$ and by hypothesis, $m > 1$. Thus, by an argument identical to that of Lemma 3.8, there exists a wave-busting arc $\gamma^* \subset \gamma_1$ for (β, γ) whose endpoints q_1 and q_2 are oppositely oriented points in $\beta_1 \cap \gamma_1$. However, in this situation, q_1 and q_2 also cobound an arc $\beta_1^* \subset \beta_1$ which meets γ_2 in a single point, so that β_1^* is a wave-busting arc for (β, γ). See Figure 26 (left). By Lemma 2.9, neither $G_\beta(\gamma)$ nor $G_\gamma(\beta)$ contains a wave, and thus (β, γ) is not a Heegaard diagram for S^3 by Theorem 2.6. □

We will assume for the remainder of the section that $m = 1$. As in Section 5, for simplicity we will isotope β so that all inessential intersections of β and γ are based in α_1^+ and α_2^+, and we define the winding W_i of β_i relative to γ_i at α_i as above. The main difference in this case is that $\iota(\alpha_i, \gamma) = 1$, so W_i is a (signed) integer.

Lemma 6.3 *If $|W_i| > 1$ for both $i = 1, 2$, then (β, γ) is not a Heegaard splitting of S^3.*

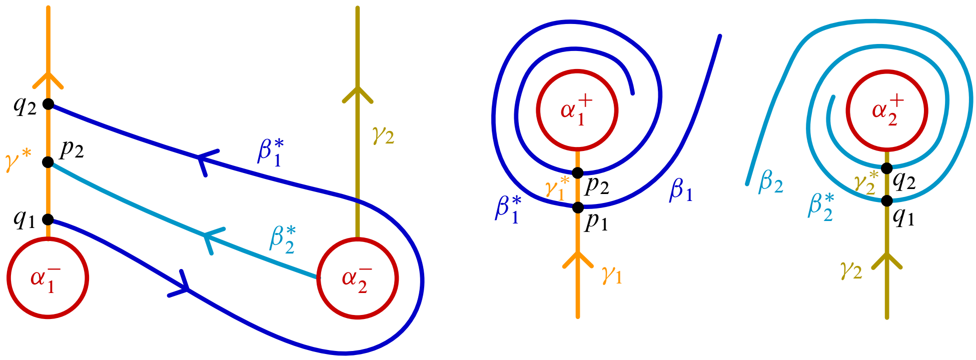

Figure 26: On the left, we see the wave-busting arcs β_1^* and γ^* that exist when $m > 1$. On the right, we see the wave-busting pairs $\{\beta_1^*, \beta_2^*\}$ and $\{\gamma_1^*, \gamma_2^*\}$ that exists when both windings are large.

Proof If $|W_i| > 1$, then there is an arc $\gamma_i^* \subset \gamma_i$ which meets β_i in its coherently oriented endpoints and avoids β in its interior. See Figure 26 (right). Thus, (γ_1^*, γ_2^*) is a wave-busting pair and $G_\beta(\gamma)$ does not contain a wave by Lemma 2.9. Let p_1, p_2 denote the endpoints of γ_1^* and q_1, q_2 denote the endpoints of γ_2^*. Note that there are arcs $\beta_1^* \subset \beta_1$ and $\beta_2^* \subset \beta_2$ which have endpoints p_1, p_2 and q_1, q_2, respectively, and which avoid γ in their interiors. Thus, (β_1^*, β_2^*) is also wave-busting pair, and we conclude that (β, γ) does not contain a wave and, as such, is not a Heegaard splitting of S^3. □

Recall from Section 3 the definition of the 0–replacement for a set of curves δ such that $\Sigma_\alpha(\delta)$ contains two arcs of slope $\frac{1}{n}$.

Lemma 6.4 *Suppose $W_2 = 0$. If β_3 and γ_3 are the 0–replacements for β and γ, respectively, then $\iota(\beta_2, \gamma_3) = \iota(\beta_3, \gamma_2) = 1$. In addition, if $W_1(\beta_1) \geq 2$, then $\iota(\beta_3, \gamma_3) = W_1(\beta_1) - 2$.*

Proof The fact that $\iota(\beta_2, \gamma_3) = \iota(\beta_3, \gamma_2) = 1$ is clear from Figure 27. Suppose $W_1 \geq 2$. Isotope β_1 so that all of the inessential intersections of β_1 and γ_1 are based at α_1^-. Then β_3, isotoped to remain disjoint from β_1, has precisely $W_1 - 2$ inessential intersections with γ_3, all of which are based at α_1^-. See Figure 27. □

Finally, we use the techniques developed up to this point to prove that if (α, β, γ) is a minimal trisection diagram, then each Heegaard diagram (α, β), (α, γ) and (β, γ) is standard, from which the main theorem follows.

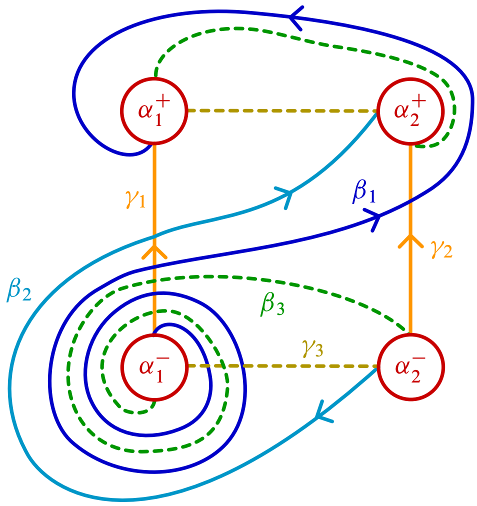

Figure 27: The 0–replacement for β and γ in the case that $W_2 = 0$ and $W_1 = 3$. The curve α_3 is suppressed for clarity; cf Figure 28.

Theorem 6.5 *Suppose that (α, β, γ) is a minimal diagram for a $(2,0)$–trisection $X = X_1 \cup X_2 \cup X_3$. Then $\iota(\alpha, \beta) = \iota(\alpha, \gamma) = \iota(\beta, \gamma) = 2$, and (α, β, γ) is homeomorphic to one of the standard diagrams pictured in Figure 2.*

Proof By Proposition 3.1, if (α, β, γ) is minimal, then $\iota(\alpha, \beta) = \iota(\alpha, \gamma) = 2$. It only remains to show that $\iota(\beta, \gamma) = 2$. We may assume that in $G_\alpha(\beta, \gamma)$, the two γ–arcs have slope $\frac{1}{0}$, and by Lemmas 6.1 and 6.2, the two β–arcs have slope $\frac{1}{n}$ for some even n. Additionally, by Lemma 6.3, we may assume without loss of generality that $|W_2| \leq 1$.

First, suppose that $|n| \geq 4$. If $n < 0$, we may apply the horizontal reflection H; hence, we assume that $n > 0$. By Lemma 2.4, we have $\beta_i \cap \gamma_i$ contains $\frac{n}{2} - 1$ essential intersections, and $\beta_i \cap \gamma_j$ contains $\frac{n}{2}$ essential intersections for $i \neq j$. Choose orientations on β and γ so that all essential intersections are coherently oriented. Since $n \geq 4$, each of β_1 and β_2 intersects both γ_1 and γ_2 essentially. This implies that $G_\beta(\gamma)$ contains an edge connecting β_1^+ to β_2^+ and $G_\gamma(\beta)$ contains an edge connecting γ_1^+ to γ_2^+. If $\beta_1 \cap \gamma_1$ contains an inessential point of intersection which is oriented opposite the essential intersections, then there is an arc $\beta_1^* \subset \beta_1$ connecting an inessential point of $\beta_1 \cap \gamma_1$ and an essential point of $\beta_1 \cap \gamma_2$. Thus, $G_\gamma(\beta)$ contains an edge connecting γ_2^+ to γ_1^-, and, by the arguments in Lemma 2.9, $G_\gamma(\beta)$ does not contain a wave. Similarly, γ_1 contains an arc connecting an inessential point of $\beta_1 \cap \gamma_1$ to an essential point of $\beta_2 \cap \gamma_1$, thus $G_\beta(\gamma)$ does not contain a wave by the same reasoning. A parallel argument shows that if $\beta_2 \cap \gamma_2$ contains an inessential intersection oriented opposite the essential intersections, then neither $G_\gamma(\beta)$ nor $G_\beta(\gamma)$ contains a wave, so by Theorem 2.6, (β,γ) is not a Heegaard diagram for S^3.

Therefore, we may assume that all points of $\beta\cap\gamma$ are coherently oriented, so that $W_i \geq 0$ for $i = 1,2$ and (β,γ) is a positive Heegaard diagram. If $W_2 = 1$, then $\iota(\beta_2,\gamma_1) = \iota(\beta_2,\gamma_2) = \frac{n}{2} > 1$. But this implies that $\frac{n}{2}$ divides $\det(M(\beta,\gamma))$, so that $M(\beta,\gamma) \neq S^3$ by Lemma 2.1, a contradiction. It follows that $W_2 = 0$.

A similar argument shows that $W_1 \neq 1$. If $W_1 \geq 2$, let a be a seam of slope $\frac{0}{1}$, let $\alpha_3 = \Delta_a$, and let β_3 and γ_3 be the 0–replacements for β and γ, respectively, as in Figure 27. In addition, let $\alpha' = \{\alpha_1,\alpha_3\}$, $\beta' = \{\beta_2,\beta_3\}$ and $\gamma' = \{\gamma_2,\gamma_3\}$, so that $\iota(\alpha',\beta') = \iota(\alpha',\gamma') = 2$. Using Lemma 6.4, we compute

$$\begin{aligned}\iota(\beta',\gamma') &= \iota(\beta_2,\gamma_2)+\iota(\beta_2,\gamma_3)+\iota(\beta_3,\gamma_2)+\iota(\beta_3,\gamma_3)\\ &= \left(\frac{n}{2}-1\right)+1+1+W_1(\beta_1)-2\\ &= \frac{n}{2}+W_1(\beta_1)-1.\end{aligned}$$

In addition, we have

$$\begin{aligned}\iota(\beta,\gamma) &= \iota(\beta_1,\gamma_1)+\iota(\beta_1,\gamma_2)+\iota(\beta_2,\gamma_1)+\iota(\beta_2,\gamma_2)\\ &= \left(\frac{n}{2}-1+W_1(\beta_1)\right)+\frac{n}{2}+\frac{n}{2}+\left(\frac{n}{2}-1\right)\\ &= 2n+W_1(\beta_1)-2.\end{aligned}$$

By the minimality of (α,β,γ), we have

$$2n+W_1(\beta_1)-2 \leq \frac{n}{2}+W_1(\beta_1)-1, \tag{3}$$

and thus $n \leq \frac{2}{3}$, a contradiction.

The only remaining possibility in this case is that $W_1 = 0$. However, this implies that the intersection matrix of (β,γ) is

$$M(\beta,\gamma) = \begin{bmatrix} \frac{n}{2}-1 & \frac{n}{2} \\ \frac{n}{2} & \frac{n}{2}-1 \end{bmatrix}.$$

By Lemma 2.1, we have $|\det(M(\beta,\gamma))| = |n-1| = 1$, and so $n = 0$ or $n = 2$, contradicting the assumption that $n \geq 4$.

Now suppose that $n = 2$. In this case, intersection matrix for (β,γ) is

$$M(\beta,\gamma) = \begin{bmatrix} W_1 & 1 \\ 1 & W_2 \end{bmatrix}.$$

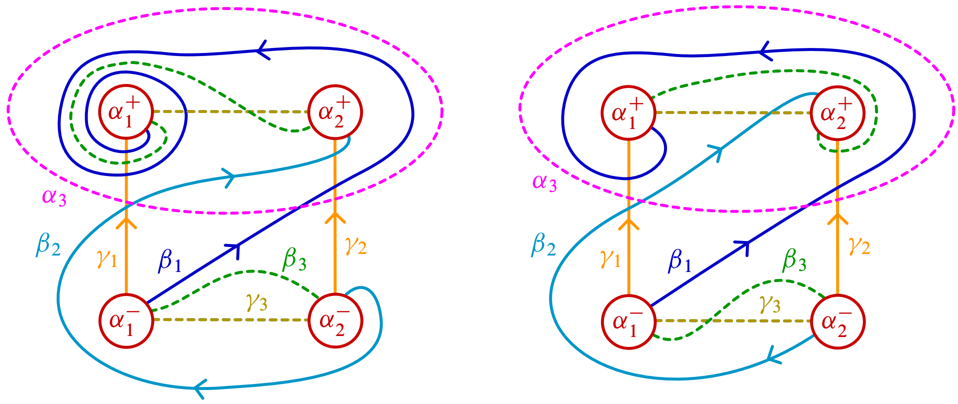

Figure 28: The 0–replacements in the cases where the slope of the β–arcs is $\frac{1}{2}$ and either $W_1 = 2$ and $W_2 = 1$ (left) or $W_1 = 1$ and $W_2 = 0$ (right)

It follows from Lemma 2.1 that $W_1 W_2 = 2$ or $W_1 W_2 = 0$. We may observe that if $W_1, W_2 \leq 0$, we may replace β and γ with curves having the same slopes in Σ_α and nonnegative winding by applying $H \circ \varphi^{-1}$ to Σ_α. See Figure 25. Thus, suppose first that $W_2 = 1$, so that $W_1 = 2$. In this case, there is a sequence of handle slides that reduces complexity. To see this, let a be a seam of slope $\frac{0}{1}$, let $\alpha_3 = \Delta_a$, and let β_3 and γ_3 be the 0–replacements for β and γ, respectively, as in Figure 28 (left). In addition, let $\alpha' = \{\alpha_1, \alpha_3\}$, $\beta' = \{\beta_2, \beta_3\}$ and $\gamma' = \{\gamma_2, \gamma_3\}$. Then, $\iota(\alpha', \beta') = \iota(\alpha', \gamma') = 2$, but $\iota(\beta', \gamma') = \iota(\beta, \gamma) - 3 = 2$. It follows that the $\mathcal{C}(\alpha', \beta'\gamma') < \mathcal{C}(\alpha, \beta, \gamma)$, a contradiction.

In the other remaining case, suppose that $W_2 = 0$. If $W_1 \geq 2$, we note that Inequality (3) is true in this case as well, and thus $n \leq \frac{2}{3}$ by above arguments, a contradiction. If $W_1 = 1$, then as above we perform handle slides in α, β and γ to reduce complexity: Let a be a seam of slope $\frac{0}{1}$, let $\alpha_3 = \Delta_a$, and let β_3 and γ_3 be the 0–replacements for β and γ, respectively, as in Figure 28 (right). In this case, the handle slides are slightly different than above; let $\alpha' = \{\alpha_2, \alpha_3\}$, $\beta' = \{\beta_1, \beta_3\}$ and $\gamma' = \{\gamma_1, \gamma_3\}$. Then $\iota(\alpha', \beta') = \iota(\alpha', \gamma') = 2$, but $\iota(\beta', \gamma') = \iota(\beta, \gamma) - 1 = 2$, contradicting that (α, β, γ) is minimal. The only remaining case when $n = 2$ is that $W_1 = W_2 = 0$; thus (α, β, γ) is the standard diagram in Figure 2 (top) (with β and γ reversed).

Finally, suppose that $n = 0$, so that both the β and the γ arcs have the same slope. In this case,

$$M(\beta, \gamma) = \begin{bmatrix} W_1 & 0 \\ 0 & W_2 \end{bmatrix},$$

and thus $W_1, W_2 = \pm 1$ by Lemma 2.1. It follows that (α, β, γ) is homeomorphic to one of the standard diagrams in Figures 2 (middle) and (bottom). (If $W_1 = W_2 = -1$, then apply H to recover Figure 2 (middle).) $\square$

7 Cosmetic surgery, trisections and exotic 4–manifolds

In the final section, we complete the proof of Theorem 1.3 by proving Proposition 7.3. We also explain and prove Corollary 1.4.

Let $L = K_1 \cup \cdots \cup K_n$ be an n–component link in a compact 3–manifold Y and let T_i denote $\partial\overline{\nu(K_i)}$ in the exterior Y_L of L, where $Y_L = Y \setminus \nu(L)$. Suppose $\vec{\rho} = (\rho_1, \dots, \rho_n)$ is a framing on L; that is, each ρ_i is a boundary slope in T_i. Let $Y_{\vec{\rho}}(L)$ denote the 3–manifold obtained by Dehn surgery on L with slope $\vec{\rho}$. We will discuss surgeries having the property that $Y_{\vec{\rho}}(L) \cong Y$. Clearly, $Y_{(1/0,\dots,1/0)}(L) = Y$, and so this case is completely uninteresting. However, surgeries with the property $Y_{\vec{\rho}}(L) \cong Y$ and $\rho_i \neq \frac{1}{0}$ for all i are quite uncommon, and we call such a surgery *cosmetic*.

Now, suppose that X is a closed 4–manifold admitting a (g,k)–trisection. Then X has a handle decomposition consisting of one 0–handle, k 1–handles, $g-k$ 2–handles, k 3–handles and one 4–handle. Let $X^{(i)}$ denote the union of the handles of index at most i. By [14], there is a unique way to attach the union of 3–handles and the 4–handle to the boundary of $X^{(2)}$, so $\partial X^{(1)} \cong \partial X^{(2)} \cong \#^k(S^1 \times S^2)$. Let L denote the union of the attaching circles of the 2–handles. Then $\partial X^{(2)}$ is obtained from surgery on L in $\partial X^{(1)}$, and thus L is a $(g-k)$–component link in $\#^k(S^1 \times S^2)$ with a cosmetic surgery.

We use this set-up to conclude our analysis of genus-two trisections.

7.1 $(2,1)$–trisections are standard

If X admits a $(2,1)$–trisection, then X has a decomposition with one 1–handle, one 2–handle and one 3–handle. The 2–handle, call it h, is attached along a knot L in $S^1 \times S^2$, the framing of h gives a cosmetic surgery slope for L in $S^1 \times S^2$. However, by work of Gabai, $S^1 \times S^2$ admits no nontrivial cosmetic surgeries [5]. More specifically, we have the following theorem.

Theorem 7.1 [5] *Suppose L is a knot in $S^2 \times S^1$ with a cosmetic surgery. Then L is a (± 1)–framed unknot.*

A Heegaard splitting of a compact 3–manifold Y with boundary is a decomposition of Y as $C_1 \cup_\Sigma C_2$, where C_1 and C_2 are compression bodies. We will require the following generalization of Haken's lemma.

Lemma 7.2 [10] *Let Y be a compact 3–manifold with a Heegaard splitting $Y = C_1 \cup_\Sigma C_2$, and suppose that D is a properly embedded disk in Y. Then there exists a disk $D' \subset Y$ such that $\partial D' = \partial D$ and $D' \cap \Sigma$ is a single simple closed curve.*

For two distinct trisections (of arbitrary genus) $X' = X_1' \cup X_2' \cup X_3'$ and $X'' = X_1'' \cup X_2'' \cup X_3''$, there is a natural trisection of the connected sum $X' \# X''$. The easiest way to see this is to take the connected sum of trisection diagrams $(\Sigma', \alpha', \beta', \gamma')$ and $(\Sigma'', \alpha'', \beta'', \gamma'')$, which yields an induced trisection diagram

$$(\Sigma' \# \Sigma'', \alpha' \cup \alpha'', \beta' \cup \beta'', \gamma' \cup \gamma'')$$

for the natural trisection of $X' \# X''$. Note that the induced diagram has the property that there is a separating curve in $\Sigma = \Sigma' \# \Sigma''$ which bounds a compressing disk in each of the three 3–dimensional handlebodies determined by the three sets of attaching curves. In this case we call the trisection *reducible*. Conversely, if $X = X_1 \cup X_2 \cup X_3$ is a reducible trisection, then it can be written as the connected sum of trisections of 4–manifolds X' and X'' such that $X = X' \# X''$. For more details, see [6].

In the following proposition, we show that every $(2,1)$–trisection is reducible, and, as such, can be written as the connected sum of a $(1,0)$– and a $(1,1)$–trisection. Since the only X' admitting a $(1,0)$–trisection are $\mathbb{CP}^2$ and $\overline{\mathbb{CP}}^2$, the only X'' admitting a $(1,1)$–trisection is $S^1 \times S^3$, and the genus-one trisections of these manifolds are unique up to diffeomorphism [6], this completes the classification of $(2,1)$–trisections.

Proposition 7.3 *Every $(2,1)$–trisection is reducible.*

Proof Let $X = X_1 \cup X_2 \cup X_3$ be a $(2,1)$–trisection with trisection surface Σ. As described above, X admits a handle decomposition with one 1–handle, one 2–handle which we will call h and one 3–handle. Let $Y = \partial X_1 = \partial X^{(1)} \cong S^2 \times S^1$. By Lemma 13 of [6], there is a trisection diagram (α, β, γ) for the trisection so that (α, β) is a Heegaard diagram for $Y \cong S^2 \times S^1$, the Heegaard diagram (α, γ) satisfies $\iota(\alpha_1, \gamma_1) = 1$ and $\alpha_2 = \gamma_2$, and h is attached to Y along γ_1 with framing given by the surface framing of γ_1 in Σ.

Let $H_{12} = X_1 \cap X_2$ and $H_{13} = X_1 \cap X_3$ be the 3–dimensional handlebodies determined by α and β, respectively, so that $H_{23} = X_2 \cap X_3$ is the handlebody determined by γ. We will show that there is a separating curve $\delta \subset \Sigma$ which bounds a disk in each of H_{12}, H_{13} and H_{23}. By Theorem 7.1, γ_1 bounds a disk in $Y = H_{12} \cup_\Sigma H_{13}$, and γ_1 has surface framing ± 1 in Σ. Since $\iota(\alpha_1, \gamma_1) = 1$ and α_1 bounds a disk in H_{12}, a positive or negative Dehn twist of γ_1 about α_1 yields a curve β_* which has surface framing zero and is isotopic in H_{12} to γ_1. Note that β_* is also isotopic to a core of H_{12} and that $\iota(\alpha_1, \beta_*) = \iota(\gamma_1, \beta_*) = 1$.

Since β_* is a 0–framed unknot in Σ, there is an embedded disk $D \subset Y$ such that $\partial D = \beta_*$ and a collar neighborhood of ∂D is disjoint from Σ. Let D' be the image of D under an isotopy which pushes β_* slightly into the interior of H_{12}, and let $\beta' = \partial D'$.

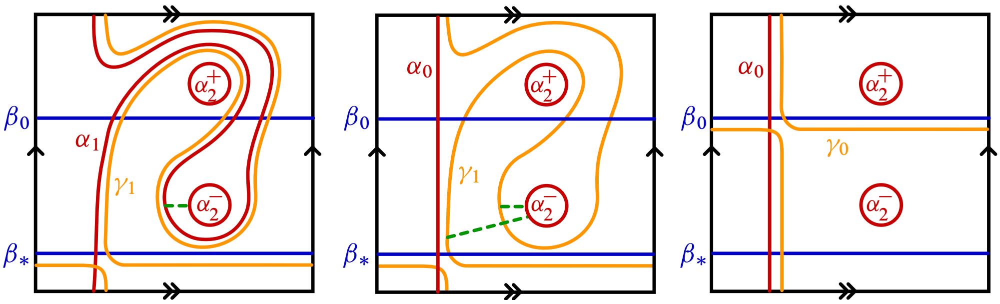

Figure 29: Three pictures of $T_1 \cup (1\text{–handle})$ showing wave moves taking α_1 and γ_1 to α_0 and γ_0

Then β' is a core of H_{12} and $C = H_{12} \setminus \nu(\beta')$ is a compression body. In addition, let $T = \partial\overline{\nu(\beta')}$, so that $D \setminus \nu(\beta')$ is a compressing disk for $Y \setminus \nu(\beta')$ with boundary in T. Then $Y \setminus \nu(\beta') = C \cup_\Sigma H_{13}$ is a Heegaard splitting satisfying the hypotheses of Lemma 7.2. It follows that there is a disk D_0 properly embedded in $Y \setminus \nu(\beta')$ such that $\partial D_0 = D' \cap T$ and such that D_0 intersects Σ in a single simple closed curve β_0. This implies that β_0 bounds a disk in H_{13}.

The compression body C may be viewed as the union of product neighborhood $T \times I$, where $T = T \times \{0\}$, and a three–dimensional 1–handle attached along $T_1 = T \times \{1\}$ whose cocore is a disk bounded by α_2 in H_{12}. By construction, there are annuli $A' \subset D' \cap C$ and $A_0 \subset D_0 \cap C$ such that the boundary of A' in T agrees with the boundary of A_0 in T, while $\partial A' \cap \Sigma = \beta_*$ and $\partial A_0 \cap \Sigma = \beta_0$. Moreover, since T is incompressible in C, these annuli are incompressible. By a standard innermost disk and outermost arc argument, each annulus may be isotoped into $T \times I$ so that it is vertical with respect to the product structure $T \times I$. It follows that, after isotopy, β_* and β_0 are disjoint and thus isotopic in T_1, although (depending on where the 1–handle is attached) β_* and β_0 may not be isotopic in $\Sigma = \partial_+ C$. See Figure 29 (left).

Since β_* is the result of Dehn twisting γ_1 around α_1, we have $\nu(\alpha_1 \cup \beta_*) = \nu(\alpha_1 \cup \gamma_1)$. The same may not be true for α_1 and γ_1 with respect to β_0; in fact, it is possible that $\iota(\alpha_1, \beta_0), \iota(\gamma_1, \beta_0) > 1$. See Figure 29 (left). However, there are curves α_0 and γ_0 isotopic to α_1 and γ_1 in T_1 such that $\iota(\alpha_0, \beta_0) = \iota(\alpha_0, \gamma_0) = \iota(\beta_0, \gamma_0) = 1$ and β_0 is the result of Dehn twisting γ_0 around α_0. Moreover, by virtue of this isotopy, there is a sequence of handle slides over α_2 taking α_1 to α_0. Likewise, there is a sequence of handle slides over $\gamma_2 = \alpha_2$ taking γ_1 to γ_0, as in Figure 29 (middle) and (right). We conclude that α_0, β_0 and γ_0 bound disks in the handlebodies determined by α, β and γ.

Let $\delta = \partial\overline{\nu(\alpha_0 \cup \gamma_0)}$. Since α_0 bounds a disk $\Delta \subset H_{12}$, we have that δ bounds a disk in H_{12} consisting of two copies of Δ banded together along γ_0. A similar argument shows that γ_0 bounding a disk in H_{23} implies that δ bounds a disk in H_{23} as well. Finally, we note that δ is also equal to $\partial\overline{\nu(\alpha_0 \cup \beta_0)}$. Thus, since β_0 bounds a disk in H_{13}, the separating curve δ also bounds a disk in H_{13}. It follows that the trisection is reducible, completing the proof. □

7.2 Cosmetic 3–sphere surgeries and exotic 4–manifolds

Although neither S^3 nor $S^2 \times S^1$ admits a cosmetic surgery on a knot L, the situation is quite different when L is a 2–component link, leading to the following natural question.

Question 7.4 Which 2–component links in a 3–sphere admit a nontrivial cosmetic Dehn surgery?

Suppose that $L \subset S^3$ is such a link, so that $L(\rho_1, \rho_2) \cong S^3$. If $\rho_1, \rho_2 \in \mathbb{Z}$, then there is an associated closed 4–manifold X obtained by first attaching (ρ_i)–framed 4–dimensional 2–handles to B^4 along $L \subset S^3 = \partial B^4$, and next capping off the resulting 3–sphere with a 4–dimensional 4–handle. Since X is simply connected, a theorem of Whitehead tells us that the homotopy type of X is determined by the symmetric, bilinear intersection form Q_X. See [8] for complete details. Since X is closed, Q_X is unimodular, and since $b_2(X) = 2$, there are only three choices:

$$Q_X \in \left\{ \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}, \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \right\}.$$

By Freedman and Quinn [3; 4], it follows that X is homeomorphic to either $S^2 \times S^2$, $\mathbb{CP}^2 \# \mathbb{CP}^2$ or $\mathbb{CP}^2 \# \overline{\mathbb{CP}}^2$. Thus, one way to look for an exotic simply connected 4–manifold X with $b_2(X) = 2$ is to limit the search to those X arising from a cosmetic surgery on a 2–component link $L \subset S^3$.

A 2–component link L in S^3 is *tunnel number one* if there is an embedded arc $a \subset S^3$ such that a intersects L only in its endpoints and $S^3 \setminus \nu(L \cup a)$ is a handlebody. In this case, $\Sigma = \partial\overline{\nu(L \cup a)}$ is a genus-two Heegaard surface for S^3, and for every integral slope $\vec{\rho}$, there is an isotopy of L into Σ so that $\vec{\rho}$ is the surface framing of L in Σ. In this setting, we can apply the results of Theorem 1.3 to place restrictions on such a link.

Lemma 7.5 *Let (Σ, α, β) be a genus-two Heegaard diagram for S^3, let γ be a cut system for Σ with surface framing $\vec{\rho}$, and let L be the 2–component link in S^3 determined by γ. Then $(\Sigma, \alpha, \beta, \gamma)$ is a $(2, 0)$–trisection diagram if and only if $L(\vec{\rho}) \cong S^3$.*

Proof Let $S^3 = H_\alpha \cup_\Sigma H_\beta$ be the Heegaard splitting determined by (Σ, α, β), and note that we may construct $L(\vec{\rho})$ by first gluing two 3–dimensional 2–handles to H_α along γ to get a compact 3–manifold $H_\alpha(\gamma)$ and gluing two 3–dimensional 2–handles to H_β along γ to get $H_\beta(\gamma)$, followed by attaching $H_\alpha(\gamma)$ to $H_\beta(\gamma)$ along their common 2–sphere boundary. In addition, $H_\alpha(\gamma)$ is a 3–ball if and only if (Σ, α, β) is a Heegaard diagram for S^3, and a similar statement holds for $H_\beta(\gamma)$. Finally, $L(\vec{\rho}) = H_\alpha(\gamma) \cup H_\beta(\gamma) \cong S^3$ if and only if both $H_\alpha(\gamma)$ and $H_\beta(\gamma)$ are 3–balls, completing the proof. □

Proof of Corollary 1.4 Suppose that L is tunnel number one and L admits a cosmetic Dehn surgery with integral slope $\vec{\rho}$. Then L is a 2–component link which may be isotoped into a genus-two Heegaard surface Σ for S^3 so that the surface framing of L in Σ is $\vec{\rho}$. Now $(\Sigma, \alpha, \beta, L)$ is a $(2,0)$–trisection diagram by Lemma 7.5, so by Theorem 1.3 there is a series of handle slides on α, β and L taking (α, β, L) to one of the standard diagrams pictured in Figure 2. It follows that there is a series of handle slides on L contained in the surface Σ that converts L to a (± 1)–framed unlink or a 0–framed Hopf link. □

There is much active research devoted to finding exotic simply connected 4–manifolds with small b_2. This corollary can be viewed as a lower bound of sorts on the complexity of link surgery descriptions of such manifolds. In particular, we see that such an exotic manifold cannot be obtained by surgery on a tunnel number one link.

Finally, we remark that this corollary may be compared to the generalized property R conjecture for two-component links, which states that if a 2–component link L has a 0–framed $\#^2(S^1 \times S^2)$ surgery, then there is a series of handle slides converting L to the 0–framed unlink. See [7] for further details. For both the corollary and the conjecture, the statements involve taking a link L with a specified surgery and converting it to an obvious canonical example or examples with the same surgery via handle slides. For more details, see [17].

References

[1] **J Berge**, *Embedding the exterior of one-tunnel knots and links in the 3–sphere*, unpublished manuscript

[2] **B Farb**, **D Margalit**, *A primer on mapping class groups*, Princeton Mathematical Series 49, Princeton University Press (2012) MR

[3] **M H Freedman**, *The topology of four-dimensional manifolds*, J. Differential Geom. 17 (1982) 357–453 MR

[4] **M H Freedman**, **F Quinn**, *Topology of* 4*–manifolds*, Princeton Mathematical Series 39, Princeton University Press (1990) MR

[5] **D Gabai**, *Foliations and surgery on knots*, Bull. Amer. Math. Soc. 15 (1986) 83–87 MR

[6] **D Gay**, **R Kirby**, *Trisecting* 4*–manifolds*, Geom. Topol. 20 (2016) 3097–3132 MR

[7] **R E Gompf**, **M Scharlemann**, **A Thompson**, *Fibered knots and potential counterexamples to the property 2R and slice-ribbon conjectures*, Geom. Topol. 14 (2010) 2305–2347 MR

[8] **R E Gompf**, **A I Stipsicz**, 4*–manifolds and Kirby calculus*, Graduate Studies in Mathematics 20, Amer. Math. Soc., Providence, RI (1999) MR

[9] **A Haas**, **P Susskind**, *The geometry of the hyperelliptic involution in genus two*, Proc. Amer. Math. Soc. 105 (1989) 159–165 MR

[10] **W Haken**, *Some results on surfaces in* 3*–manifolds*, from "Studies in Modern Topology" (P J Hilton, editor), Math. Assoc. Amer. (distributed by Prentice-Hall, Englewood Cliffs, NJ) (1968) 39–98 MR

[11] **T Homma**, **M Ochiai**, **M-o Takahashi**, *An algorithm for recognizing* S^3 *in* 3*–manifolds with Heegaard splittings of genus two*, Osaka J. Math. 17 (1980) 625–648 MR

[12] **K Ishihara**, *On tunnel number one links with surgeries yielding the* 3*–sphere*, Osaka J. Math. 47 (2010) 189–208 MR

[13] **K Johannson**, *Topology and combinatorics of* 3*–manifolds*, Lecture Notes in Mathematics 1599, Springer, New York (1995) MR

[14] **F Laudenbach**, **V Poénaru**, *A note on* 4*–dimensional handlebodies*, Bull. Soc. Math. France 100 (1972) 337–344 MR

[15] **H Matsuda**, **M Ozawa**, **K Shimokawa**, *On non-simple reflexive links*, J. Knot Theory Ramifications 11 (2002) 787–791 MR

[16] **E Mayrand**, *Dehn fillings on a two torus boundary components* 3*–manifold*, Osaka J. Math. 39 (2002) 779–793 MR

[17] **J Meier**, **T Schirmer**, **A Zupan**, *Classification of trisections and the generalized property R conjecture*, Proc. Amer. Math. Soc. 144 (2016) 4983–4997 MR

[18] **J Meier**, **A Zupan**, *Bridge trisections of knotted surfaces in* S^4, preprint (2015) arXiv

[19] **M Ochiai**, *Heegaard diagrams and Whitehead graphs*, Math. Sem. Notes Kobe Univ. 7 (1979) 573–591 MR

[20] **M Ochiai**, *Heegaard diagrams of* 3*–manifolds*, Trans. Amer. Math. Soc. 328 (1991) 863–879 MR

[21] **M Teragaito**, *Links with surgery yielding the* 3*–sphere*, J. Knot Theory Ramifications 11 (2002) 105–108 MR

Department of Mathematics, Indiana University, Rawles Hall 425
831 East 3rd Street, Bloomington, IN 47405, United States

Department of Mathematics, University of Texas at Austin
1 University Station C1200, Austin, TX 78712, United States

jlmeier@indiana.edu, zupan@math.utexas.edu

Proposed: Colin Rourke
Seconded: David Gabai, Ronald Stern

Received: 1 December 2014
Revised: 21 February 2016